

HP StorageWorks

Fabric OS 5.x administrator guide

Legal and notice information

© Copyright 2005 Hewlett-Packard Development Company, L.P.

© Copyright 2005 Brocade Communications Systems, Incorporated.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, and Windows are U.S. registered trademarks of Microsoft Corporation.

UNIX is a registered trademark of The Open Group.

Linux is a U.S. registered trademark of Linus Torvalds.

Java is a U.S. trademark of Sun Microsystems, Inc.

Fabric OS 5.x administrator guide

Contents

About this guide	13
Intended audience	13
Related documentation	13
HP StorageWorks Fabric OS 5.x master glossary	13
Document conventions and symbols	14
HP technical support	14
HP-authorized reseller.	15
Helpful web sites	15
1 Introducing Fabric OS CLI procedures.	17
About procedural differences	17
Scope and references	17
About the CLI	18
Help information.	18
Displaying command help	18
Additional help topics	19
2 Performing basic configuration tasks.	21
Connecting to the Command Line Interface	21
Connecting with telnet	21
Connecting through the serial port	21
Setting the IP address	22
Setting the default account passwords	22
Changing the default passwords at login	23
Setting the date and time.	24
Setting the date and time	25
Synchronizing local time with an external source	25
Setting the time zone	25
Maintaining licensed features.	26
Unlocking a licensed feature	26
Removing a licensed feature	28
Customizing the switch name	28
Customizing the switch name.	28
Customizing the chassis name	29
Changing the chassis name.	29
Disabling and enabling a switch	29
Disabling a switch	29
Enabling a switch.	29
Disabling and enabling a port	30
Disabling a port.	30
Enabling a port	30
Activating Ports on Demand	30
Activating Ports on Demand	31
Making basic connections	31
Connecting to devices	31
Connecting to other switches.	31
Working with domain IDs	32
Displaying domain IDs	32
Setting the domain ID	33
Linking through a gateway.	33
Configuring a link through a gateway	33
Checking status	34
Checking switch operation	34
Verifying HA features	34

Verifying fabric connectivity	34
Verifying device connectivity	35
Tracking and controlling switch changes	35
Enabling the TC feature	36
Displaying the status of the TC feature	36
Viewing the switch status policy threshold values	36
Setting the switch status policy threshold values	37
3 Configuring standard security features	39
Secure protocols	39
Ensuring network security	40
Configuring the telnet interface	41
Disabling telnet	41
Enabling telnet	41
Blocking listeners	42
Accessing switches and fabrics	42
Creating and maintaining user-defined accounts	43
Displaying account information	43
Creating a user-defined account	43
Deleting a user-defined account	44
Changing account parameters	44
Recovering user-defined accounts	44
Changing an account password	45
Changing the password for the current login account	45
Changing the password for a different account	45
Setting up RADIUS AAA service	45
Configuring the RADIUS server	46
Linux	47
Adding the attribute to the server	47
Creating the user	47
Enabling clients	48
Windows 2000	48
Enabling CHAP	48
Configuring users	49
Configuring the RADIUS server	49
Configuring the switch	50
Displaying the current RADIUS configuration	50
Adding a RADIUS server to the switch configuration	51
Enabling or disabling RADIUS service	51
Deleting a RADIUS server from the configuration	51
Changing a RADIUS server configuration	51
Changing the order in which RADIUS servers are contacted for service	52
Enabling and disabling local authentication	52
Configuring for the SSL protocol	52
Browser and Java support	53
Summary of SSL procedures	53
Choosing a CA	54
Generating a public/private key	54
Generating and storing a CSR	54
Obtaining certificates	55
Installing a switch certificate	55
Activating a switch certificate	56
Configuring the browser	56
Checking and installing root certificates on Internet Explorer	56
Checking and installing root certificates on Mozilla	56
Installing a root certificate to the Java Plug-in	57
Displaying and deleting certificates	57
Troubleshooting certificates	58
Configuring SNMP agent and traps	58

Setting the security level	59
Using the snmpConfig command	59
Using legacy commands for SNMPv1	62
Configuring secure file copy	67
Setting the boot PROM password	67
4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32	68
Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director.	68
Without a recovery string	69
4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32	70
Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director.	70
Recovering forgotten passwords	71
4 Maintaining configurations and firmware	73
Maintaining configurations	73
Displaying configuration settings	73
Backing up a configuration	73
Restoring a configuration	74
Restoring configurations in a FICON environment	75
Downloading configurations across a fabric	76
Printing hard copies of switch information.	76
Maintaining firmware	76
Obtaining and unzipping firmware	76
Checking connected switches	77
About the download process.	78
Effects of firmware changes on accounts and passwords	79
Considerations for downgrading firmware	79
Considerations for FICON CUP environments.	79
Upgrading HP StorageWorks switches	79
Summary of the upgrade process	80
Upgrading 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32	80
Upgrading HP StorageWorks directors.	82
Summary of the upgrade process	82
Upgrading the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director.	83
Troubleshooting firmware downloads	85
5 Configuring Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director	87
Identifying ports	87
By slot and port number	87
By port area ID	88
Basic blade management	88
Powering port blades off and on	88
Powering off a port blade	88
Providing power to a port blade	88
Disabling and enabling port blades	89
Disabling a port blade	89
Enabling a port blade.	89
Conserving power	89
Blade terminology and compatibility	89
CP blades.	90
Port blade compatibility	91
Setting chassis configurations.	91
Obtaining slot information	92
Displaying the status of all slots in the chassis	92
Configuring a new SAN Director 2/128 with two domains	93
Converting an installed SAN Director 2/128 to support two domains	94
Setting the blade beacon mode	95

6 Routing traffic	97
About data routing and routing policies	97
Specifying the routing policy	97
Assigning a static route	98
Specifying frame order delivery	98
Forcing in-order frame delivery across topology changes	99
Restoring out-of-order frame delivery across topology changes	99
Using DLS	99
Checking and setting DLS	99
Viewing routing path information	100
Viewing routing information along a path	102
7 Administering FICON fabrics	105
FICON overview	105
Configuring switches	108
Preparing a switch	108
Configuring a single switch	108
Configuring a high-integrity fabric	109
Setting a unique domain ID	109
Displaying information	110
Link incidents	110
Registered listeners	110
Node identification data	110
FRU failures	111
Swapping ports	111
Clearing the FICON management database	111
Using FICON CUP	111
Setup summary	112
Enabling and disabling FMS mode	112
Setting up CUP when FMS mode is enabled	113
Displaying the fmsmode setting	113
Displaying mode register bit settings	113
Setting mode register bits	114
Persistently enabling and disabling ports	115
Port and switch naming standards	116
Adding and removing FICON CUP licenses	116
Zoning and PDCM considerations	116
Backing up and restoring configurations	116
Troubleshooting	116
Identifying ports	117
Backing up FICON files	118
Uploading the configuration files	118
Downloading configuration files with Active=Saved mode enabled	118
Downloading configuration files with Active=Saved mode disabled	118
Recording configuration information	118
Sample IOCP configuration file for the SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128	120
Sample Resource Management Facility configuration file for mainframe	121
8 Configuring the Distributed Management Server	123
Enabling and disabling the platform services	123
Enabling platform services	123
Disabling platform services	123
Controlling access	124
Displaying the management server ACL	124
Adding a member to the ACL	124
Deleting a member from the ACL	125
Configuring the server database	126
Viewing the contents of the management server database	126

Clearing the management server database	127
Controlling topology discovery	127
Displaying topology discovery status	127
Enabling topology discovery	127
Disabling topology discovery	128
9 Working with diagnostic features	129
Viewing POST	129
Viewing switch status	130
Viewing the overall status of the switch	130
Displaying switch information	131
Displaying the uptime for a switch	131
Viewing port information	132
Viewing the status of a port	132
Displaying the port statistics	133
Displaying a summary of port errors for a switch	133
Viewing equipment status	135
Displaying the status of the fans	135
Displaying the status of a power supply	135
Displaying temperature status	136
Viewing the system message log	136
Displaying the system message log, with no page breaks	136
Displaying the system message log, one page at a time	136
Clearing the system message log	136
Viewing the port log	137
Configuring for syslogd	138
Configuring the host	138
Configuring the switch	139
Specifying syslogd hosts	139
Setting the facility level	139
Removing a syslogd host from the list	139
Viewing and saving diagnostic information	140
Setting up automatic trace dump transfers	140
Specifying a remote server	140
Enabling the automatic transfer of trace dumps	140
Setting up periodic checking of the remote server	141
Saving a comprehensive set of diagnostic files to the server	141
10 Troubleshooting	143
Most common problem areas	143
Gathering information for technical support	144
Troubleshooting questions	144
Analyzing connection problems	145
Checking the logical connection	145
Checking for Fibre Channel connectivity problems	145
Checking the Simple Name Server (SNS)	147
Checking for zoning problems	148
Restoring a segmented fabric	148
Reconciling fabric parameters individually	148
Downloading a correct configuration	149
Reconciling a domain ID conflict	149
Correcting zoning setup issues	149
Correcting a fabric merge problem quickly	150
Verifying a fabric merge problem	150
Editing zone configuration members	151
Reordering the zone member list	151
Recognizing MQ errors	151
Correcting I2C bus errors	151
Checking fan components	152

Checking the switch temperature	152
Checking the power supply	152
Checking the temperature, fan, and power supply	152
Correcting device login issues	152
Identifying media-related issues	155
Testing a port's external transmit and receive path	156
Testing a switch's internal components	156
Testing components to and from the HBA	157
Correcting link failures	157
Determining whether the negotiation was successfully completed	157
Checking for a loop initialization failure	158
Checking for a point-to-point initialization failure	158
Correcting a port that came up in the wrong mode	159
Correcting marginal links.	159
Inaccurate information in the system message log	161
Port initialization and FCP auto-discovery process.	161
11 Administering extended fabrics	163
About extended link buffer allocation	163
SAN Switch 2/8V, SAN Switch 2/16V, and SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director (FC2-16 port blades).	163
Brocade 4Gb SAN Switch for HP p-Class BladeSystem, SAN Switch 4/32, and 4/256 SAN Director (FC4-16 and FC4-32 port blades)	163
Fabric considerations	163
Choosing an extended ISL mode	164
Configuring an extended ISL	165
Trunking over distance	166
12 Administering ISL trunking	167
About ISL trunking	167
Standard trunking criteria	168
Fabric considerations	168
Initializing trunking on ports.	169
Disabling and reenabling the switch.	169
Disabling and reenabling ports	169
Monitoring traffic	169
Using the portperfshow command	170
Enabling and disabling ISL trunking	171
Enabling or disabling ISL trunking on one port.	171
Enabling or disabling ISL trunking for all of the ports on a switch	171
Setting port speeds	171
Setting the speed for one port	172
Setting the speed for all of the ports on the switch	172
Displaying trunking information	173
Trunking over extended fabrics.	173
Troubleshooting trunking problems	174
Listing link characteristics	174
Recognizing buffer underallocation	174
Getting out of buffer-limited mode on E_Ports or LD_Ports:	174
13 Administering advanced zoning	177
Zoning terminology.	177
Zoning concepts.	178
Zone types	178
Zone objects	179
Zone aliases	180
Zone configurations	180
Zoning enforcement	181
Software-enforced zoning	181
Hardware-enforced zoning	181

Rules for configuring zones	184
Creating and managing zone aliases	185
Creating an alias	185
Adding members to an alias	186
Removing members from an alias.	186
Deleting an alias	186
Viewing an alias in the defined configuration	187
Creating and maintaining zones.	187
Creating a zone	187
Adding devices (members) to a zone	187
Removing devices (members) from a zone.	188
Deleting a zone	188
Viewing a zone in the defined configuration	188
Merging zones	188
Creating and modifying zoning configurations.	190
Creating a zoning configuration	191
Adding zones (members) to a zoning configuration	191
Removing zones (members) from a zone configuration	191
Deleting a zone configuration	191
Clearing changes to a configuration.	192
Viewing all zone configuration information	192
Viewing selected zone configuration information	192
Viewing a configuration in the effective zone database	193
Maintaining zone objects	193
Copying a zone object	193
Deleting a zone object	193
Renaming a zone object	194
Managing zoning configurations in a fabric	195
Adding a new switch or fabric	195
Splitting a fabric	197
Using zoning to administer security.	197
Resolving zone conflicts.	197
14Administering advanced performance monitoring.	199
Displaying and clearing the CRC error count	200
Monitoring EE performance	201
Adding EE monitors	201
Monitoring the traffic from Host A to Dev B	202
Monitoring the traffic from Dev B to Host A	202
Setting a mask for EE monitors.	202
Displaying the current EE mask of a port	203
Displaying a monitor	204
Deleting EE monitors	204
Monitoring filter-based performance	204
Adding standard filter-based monitors	205
Adding custom filter-based monitors	205
Adding filter-based monitors	206
Deleting filter-based monitors.	206
Monitoring ISL performance	207
Monitoring trunks	207
Displaying monitor counters	207
Clearing monitor counters	209
Saving and restoring monitor configurations	211
Collecting performance data	211
A Configuring the PID format	213
About PIDs and PID binding.	213
Summary of PID formats	213
Impact of changing the fabric PID format.	213


Host reboots	214
Static PID mapping errors	214
Changes to configuration data	214
Selecting a PID format	215
Evaluating the fabric	216
Planning the update procedure	218
Online update	218
Offline update	218
Hybrid update	219
Changing to Core PID format	219
Changing to Extended Edge PID format	220
Converting port number to area ID	222
PID format changes	224
Executing the basic procedure	224
Executing the HP-UX procedure	225
Executing the AIX procedure	227
Swapping port area IDs	227
B Configuring interoperability mode	229
C Using the HP Remote Switch feature	231
D Understanding legacy password behavior	233
Password management information	233
Password prompting behaviors	234
Password migration during firmware changes	235
Password recovery options	236
E Zone merging scenarios	237
F Upgrading firmware in single-CP mode	239
Upgrading HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32	239
Upgrading a single Core Switch 2/64 or SAN Director 2/128 blade	240
Index	243
Figures	
1 Cascaded configuration with two switches	109
2 Cascaded configuration with three switches	109
3 FICON switch configuration worksheet	119
4 Distribution of traffic over ISL trunking groups	167
5 Zoning example	177
6 Hardware-enforced non-overlap ping zones	183
7 Hardware-enforced overlapping zones	183
8 Zoning with hardware assist (mixed-port and WWN zones)	184
9 Session-based hard zoning	184
10 Setting EE monitors on a port	202
11 Proper placement of EE performance monitors	202
12 Mask positions for EE monitors	203
13 4/256 SAN Director with Extended Edge PID	223
Tables	
1 Document conventions	14
2 Maximum number of simultaneous sessions	22
3 Conversion from UTC to local time	26
4 Secure protocol support	39
5 Items needed to deploy secure protocols	39
6 Main security scenarios	40
7 Blocked listener applications	42
8 Access defaults	42

9	SSL certificate files	53
10	Commands to display and delete SSL certificates	57
11	SSL messages and actions	58
12	Backup and restore in a FICON CUP environment	75
13	Recommended firmware	78
14	Effects of firmware changes on accounts and passwords	79
15	HP StorageWorks director terminology and abbreviations	90
16	Blades supported by each HP StorageWorks director	91
17	Supported configuration options	91
18	Fabric OS commands related to FICON and FICON CUP	107
19	FICON CUP mode register bits	114
20	Port error summary description	134
21	Commands for port log management	137
22	Fabric OS and UNIX message severities	138
23	Common troubleshooting problems and tools	143
24	Types of zone discrepancies	149
25	Commands for debugging zoning	150
26	Component test descriptions	156
27	Switch component tests	157
28	Switchshow output and suggested action	159
29	Loopback modes	160
30	Extended ISL modes: switches with Bloom ASIC	164
31	Extended ISL modes: switches with Goldeneye ASIC	164
32	Extended ISL modes: switches with Condor ASIC	165
33	Types of zoning	178
34	Approaches to fabric-based zoning	179
35	Enforcing hardware zoning	182
36	Resulting database size: 0 to 96K	189
37	Resulting database size: 96K to 128K	189
38	Resulting database size: 128K to 256K	190
39	Resulting database size: 256K to 1M	190
40	Zoning database limitations	195
41	Considerations for zoning architecture	198
42	Advanced performance monitoring commands	199
43	Commands to add filter-based monitors	205
44	Predefined values at offset 0	206
45	Effects of PID format changes on configurations	215
46	PID format recommendations for adding new switches	216
47	Earliest Fabric OS versions for Extended Edge PID format	220
48	Account and password characteristics matrix	233
49	Password prompting matrix	234
50	Password migration behavior during firmware upgrade and downgrade	235
51	Password recovery options	236
52	Zone merging scenarios	237

About this guide

This guide provides information about:

- Fabric OS procedures
- Basic configuration tasks
- Security features
- Diagnostics
- Extended fabrics
- ISL trunking
- Zoning
- Performance monitoring

 **NOTE:** FICON is not supported on HP B-Series Fibre Channel switches. The FICON information in this document is included for reference only.

Intended audience

This guide is intended for:

- System administrators responsible for setting up HP StorageWorks Fibre Channel Storage Area Network (SAN) switches
- Technicians responsible for maintaining the Fabric Operating System (OS)

Related documentation

Documentation, including white papers and best practices documents, is available on the HP web site:

<http://www.hp.com/country/us/eng/prodserv/storage.html>.

IMPORTANT: For late breaking, supplemental information, access the latest version of the *HP StorageWorks Fabric OS 5.x release notes* using the following steps.

To access current Fabric OS related documents:

1. Locate the **IT storage products** section of the web page.
2. Under **Networked storage**, click **SAN infrastructure**.
3. From the **SAN Infrastructure** web page, locate the **SAN Infrastructure products** section.
4. Click **Fibre Channel Switches**.
5. Locate the **B-Series Fabric-Enterprise Class** section. Click **4/256 SAN Director and 4/256 SAN Director power pack**, to access Fabric OS 5.x documents (such as this document).
The switch overview page displays.
6. Go to the **Product Information section**, located on the right side of the web page.
7. Click **Technical documents**.
8. Follow the onscreen instructions to download the applicable documents.

HP StorageWorks Fabric OS 5.x master glossary

This guide uses industry standard SAN terminology. However, some terms are intrinsic to Fabric OS 5.x. See the *HP StorageWorks Fabric OS 5.x master glossary* for a complete list of terms and definitions.

Access the master glossary from the HP StorageWorks SAN Switch Documentation CD that shipped with your switch. Also, access from the HP web site using the procedure outlined in "[Related documentation](#)".


Document conventions and symbols


Table 1 Document conventions

Convention	Element
Medium blue text: Figure 1	Cross-reference links and e-mail addresses
Medium blue, underlined text (http://www.hp.com)	Web site addresses
Bold font	<ul style="list-style-type: none">• Key names• Text typed into a GUI element, such as into a box• GUI elements that are clicked or selected, such as menu and list items, buttons, and check boxes
<i>Italics font</i>	Text emphasis
Monospace font	<ul style="list-style-type: none">• File and directory names• System output• Code• Text typed at the command-line
<i>Monospace, italic font</i>	<ul style="list-style-type: none">• Code variables• Command line variables
Monospace, bold font	Emphasis of file and directory names, system output, code, and text typed at the command line

 **WARNING!** Indicates that failure to follow directions could result in bodily harm or death.

 **CAUTION:** Indicates that failure to follow directions could result in damage to equipment or data.

 **IMPORTANT:** Provides clarifying information or specific instructions.

 **NOTE:** Provides additional information.

 **TIP:** Provides helpful hints and shortcuts.

HP technical support

Telephone numbers for worldwide technical support are listed on the HP support web site:
<http://www.hp.com/support/>.

Collect the following information before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages

- Operating system type and revision level
- Detailed, specific questions

For continuous quality improvement, calls may be recorded or monitored.

HP strongly recommends that customers sign up online using the Subscriber's choice web site:

<http://www.hp.com/go/e-updates>.

- Subscribing to this service provides you with e-mail updates on the latest product enhancements, newest versions of drivers, and firmware documentation updates as well as instant access to numerous other product resources.
- After signing up, you can quickly locate your products by selecting **Business support** and then **Storage** under Product Category.

HP-authorized reseller

For the name of your nearest HP-authorized reseller:

- In the United States, call 1-800-282-6672.
- Elsewhere, visit the HP web site: <http://www.hp.com>. Click **Contact HP** to find locations and telephone numbers.

Helpful web sites

For other product information, see the following HP web sites:

- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- <http://www.hp.com/support/>
- <http://www.docs.hp.com>

1 Introducing Fabric OS CLI procedures

This chapter contains procedures for configuring and managing an HP StorageWorks Storage Area Network (SAN) using the Fabric OS Command Line Interface (CLI).

The guide applies to the following HP StorageWorks product models:

- HP StorageWorks switches: 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32

These HP StorageWorks models contain a fixed number of ports (they are *fixed-port switches*). The SAN Switch 4/32, 4/8 SAN Switch, and 4/16 SAN Switch allow you to license and activate extra fixed ports with the Ports on Demand feature.

- HP StorageWorks directors: Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director

These HP StorageWorks models can contain a variable number of ports, which you install by plugging port blades into the director chassis. The 4/256 SAN Director can have up to 256 ports; the Core Switch 2/64 and SAN Director 2/128 can have up to 128 ports.

About procedural differences

As a result of the differences between fixed-port and variable-port devices, procedures sometimes differ among HP StorageWorks models. Also, because the domain architecture of the Core Switch 2/64 differs from that of the SAN Director 2/128 and 4/256 SAN Director, there are sometimes procedural differences among these models. As new HP StorageWorks models are introduced, new features sometimes apply only to those models.

When procedures or parts of procedures apply to some models but not others, this guide identifies the specifics for each model. For example, a number of procedures that apply only to variable-port devices are found in "[Configuring Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director](#)" on page 87. Procedures that apply only to the SAN Switch 4/32 are labeled as such.

 **NOTE:** When command examples in this guide show user input enclosed in quotation marks, the quotation marks are required for versions of Fabric OS earlier than 4.0.0. They are optional in later versions, unless specifically called for in the procedures.

Scope and references

Although many different software and hardware configurations are tested and supported by HP, documenting all possible configurations and scenarios is beyond the scope of this guide. In some cases, earlier releases of Fabric OS are documented to present considerations for interoperating with them.

The installation guides for HP StorageWorks products describe how to power up devices and set their IP addresses. After the IP address is set, you can use the CLI procedures contained in this guide.

This guide provides only the level of detail required to perform the procedures. If you need more information about the commands used in the procedures, see online help or the *HP StorageWorks Fabric OS 5.x command reference guide*.

You can use several access methods to configure a switch:

- CLI
 - A telnet session into logical switches
 - A telnet session into active and standby CPs for director class switches
 - A serial console, including active and standby CPs for director class switches
 - An optional modem, which behaves like a serial console port

For CLI details, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

- Advanced Web Tools: For Advanced Web Tools procedures, see the *HP StorageWorks Fabric OS 5.x Advanced Web Tools administrator guide*.
- Fabric Manager: For Fabric Manager procedures, see the *HP StorageWorks Fabric Manager 5.x administrator guide*.
- A third-party application using the API: For third-party application procedures, see the third-party API documentation.

About the CLI

The Fabric OS CLI is the complete fabric management tool for HP StorageWorks SANs that enables you to:

- Access the full range of Fabric OS features, based on license keys
- Configure, monitor, dynamically provision, and manage every aspect of the SAN
- Configure and manage the HP StorageWorks fabric on multiple, efficient levels
- Identify, isolate, and manage SAN events across every switch in the fabric
- Manage switch licenses
- Perform fabric stamping

To manage a switch using telnet, Simple Network Management Protocol (SNMP), and Advanced Web Tools, the switch must be connected to a network through the switch Ethernet port (out of band) or from the Fibre Channel (in band). The switch must be configured with an IP address to allow for the network connection. See the installation guide for your switch model for information on physically connecting to the switch.

You can access switches from different connections, such as Advanced Web Tools, CLI, and API. When these connections are simultaneous, changes from one connection might not be updated to the other, and some modifications might be lost. When simultaneous connections are used, make sure that you do not overwrite the work of another connection.

In a mixed fabric containing switches running various Fabric OS versions, you should use the latest-model switches running the most recent release for the primary management tasks. The principal management access should be set to the core switches in the fabric. For example, to run Secure Fabric OS, use the latest-model switch as the primary Fabric Configuration Server (FCS), the location to perform zoning tasks, and the time server.

A number of management tasks are designed to make fabric-level changes; for example, zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent tasks. For a large fabric, it might take a few minutes.

Help information

Each Fabric OS command provides Help information that explains the command function, its possible operands, its level in the command hierarchy, and additional pertinent information.

Displaying command help

To display help information:

1. Connect to the switch and log in as admin.
2. To display a list of all command help topics for a given login level, issue the `help` command with no arguments.

For example, if you are logged in as user and issue the `help` command, a list of all user-level commands that can be executed is displayed. The same rule applies to the admin role. In addition, any user-configured command that uses a switchAdmin role also displays commands available to users with the switchAdmin role.
3. To display help for a specific command for a given login level, issue `help command`, where *command* is the name of the command for which you need information.

For example:

```
switch:admin> help configure
Administrative Commands                                configure(1m)
NAME
    configure - change system configuration settings
SYNOPSIS
    configure
AVAILABILITY
    admin
DESCRIPTION
    This command changes some system configuration settings,
    including:
    o Arbitrated loop settings
    o Switch fabric settings
    o System services settings
    o Virtual channel settings
    (output truncated)
```

Additional help topics

The following commands provide help files for specific topics:

- `diagHelp` provides diagnostic information
- `ficonHelp` provides Fibre Connection (FICON) information
- `fwHelp` provides Fabric Watch information
- `licenseHelp` provides license information
- `perfHelp` provides Performance Monitoring information
- `routeHelp` provides routing information
- `trackChangesHelp` provides Track Changes (TC) information
- `zoneHelp` provides zoning information

2 Performing basic configuration tasks

This chapter contains procedures for performing basic switch configuration tasks using the Fabric OS CLI.

Connecting to the Command Line Interface

You can connect to the CLI either through a telnet connection or through the serial port.

Connecting with telnet

1. Verify that the switch is connected to the IP network through the RJ-45 Ethernet port.
Switches in the fabric that are not connected through Ethernet can be managed through switches that use IP over Fibre Channel. The embedded port must have an assigned IP address.
2. Open a telnet connection using the IP address of the logical switch to which you want to connect.
If you telnet to the active Control Processor (CP) or log in to the active CP console, you are prompted for the switch number when the platform is set up in dual (or multiple) switch mode. For example, the SAN Director 2/128 does not prompt you if you are using configuration option 1, but does prompt you if you have used configuration options 2–4. See ["Configuring Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director"](#) on page 87 for details about director configuration options.
The Core Switch 2/64 and SAN Director 2/128 (configured with two domains) have two logical switches (sw0 and sw1).
The login prompt is displayed when the telnet connection finds the switch in the network.
3. Enter the account ID (defaults are `user` or `admin`) at the login prompt.
4. Enter the password.
The default password is `password`.
If you have not changed the system passwords from the default, you are prompted to change them. Enter the new system passwords, or press **Ctrl-c** to skip the password prompts.
5. Verify that the login was successful.
The prompt displays the switch name and user ID to which you are connected.

```
login: admin
password: xxxxxxxx
switch:admin>
```

6. Observe the following considerations for telnet connections:
 - Never change the IP address of the switch while two telnet sessions are active; if you do, your next attempt to log in fails. To recover, gain access to the switch by one of these methods:
 - Perform a fast boot using Advanced Web Tools. When the switch comes up, the telnet quota is cleared. (For instructions on performing a fast boot with Advanced Web Tools, see the *HP StorageWorks Fabric OS 5.x Advanced Web Tools administrator guide*.)
 - If you have the required privileges, connect through the serial port, log in as root, and use operating system commands to identify and kill the telnet processes without disrupting the fabric.
 - For admin level accounts, Fabric OS limits the number of simultaneous telnet sessions per switch to two. For details on session limits, see ["Configuring the telnet interface"](#) on page 41 and ["Creating and maintaining user-defined accounts"](#) on page 43.

Connecting through the serial port

1. Connect the serial cable to the serial port on the switch and to an RS-232 serial port on the workstation.
If the serial port on the workstation is RJ-45 instead of RS-232, remove the adapter on the end of the serial cable and insert the exposed RJ-45 connector into the RJ-45 serial port on the workstation.

2. Open a terminal emulator application (such as HyperTerminal on a PC, or TERM, TIP, or Kermit in a UNIX® environment), and configure the application as follows:

- In a Windows® environment:

<u>Parameter</u>	<u>Value</u>
Bits per second	9600
Databits	8
Parity	None
Stop bits	1
Flow control	None

- In a UNIX environment, enter the following string at the prompt:

```
tip /dev/ttyb -9600
```

If ttyb is already in use, you can use ttya (enter `tip /dev/ttya -9600`).

3. Observe the following considerations for serial connections:

- Some procedures require that you connect through the serial port, for example, setting the IP address or setting the boot PROM password.
- If secure mode is enabled, connect through the serial port of the primary FCS switch.
- For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, you can connect to CP0 or CP1 using either of the two serial ports.

Setting the IP address

You must connect through the serial port to set the IP address (see ["Connecting through the serial port"](#) on page 21). After connecting, issue the `ipAddrSet` command to set the IP address.

△ **CAUTION:** The use of IP address 0.0.0.0 is not supported. Do not use this address.

Fabric OS versions beginning with 2.6.0, 3.1.0, and 4.0.0 support Classless Inter-Domain Routing (CIDR).

Setting the default account passwords

For each logical switch (domain), there are admin and user default access accounts. These accounts designate the following levels of authorization—called *roles*—for using the system:

- Admin level for administrative use
- User level for non-administrative use, such as monitoring system activity
- SwitchAdmin level for administrative use, except for security, user management, and zoning

Two accounts—factory and root—are reserved for development and manufacturing. You can change their passwords, which is optional, but do not use these accounts under normal circumstances.

Table 2 shows the number of simultaneous login sessions allowed for each role.

Table 2 Maximum number of simultaneous sessions

User name	Maximum sessions
admin	2
user	4

For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, SAN Switch 4/32, SAN Director 2/128, and the 4/256 SAN Director (default configuration with one domain), there is one set of default access accounts.

For the Core Switch 2/64 and SAN Director 2/128 (configured with two domains), each logical switch has its own set of default access accounts. The default account names and passwords are the same for both of the logical switches.

You can also create up to 15 additional accounts per logical switch and designate their roles as either admin, switchAdmin, or user. See the procedures for doing so in ["Creating and maintaining user-defined accounts"](#) on page 43.

For large enterprises, Fabric OS supports RADIUS services, as described in ["Setting up RADIUS AAA service"](#) on page 45.

In addition to the account access passwords, each switch can set a boot PROM password. For greater security, HP recommends that you set this password to protect system boot parameters from unauthorized access. See ["Setting the boot PROM password"](#) on page 67.

Each of the default access accounts has an associated password. The first time you connect to a Fabric OS switch, you are prompted to change these default account passwords.

If you do not change the default passwords, you are prompted to do so at each subsequent login until all system passwords have been changed from their default values. Thereafter, use the `passwd` command to change passwords.

For more background information on passwords, see ["Changing an account password"](#) on page 45.

Changing the default passwords at login

1. Connect to the switch and log in as admin.


The default password for all default accounts is `password`.

2. At each of the `Enter new password` prompts, either enter a new password or skip the prompt.

You can skip a prompt by pressing **Enter**. You can bypass all further prompts by pressing **Ctrl-c**.

Although the root and factory accounts are not meant for general use, change their passwords if prompted to do so, and save the passwords in case they are needed for recovery purposes.


You cannot reuse the default passwords.

 **NOTE:** Record the passwords exactly as entered and store them in a secure place; recovering passwords requires significant effort and fabric downtime. The initial login prompt accepts a maximum password length of eight characters. Characters beyond the eighth are ignored. Only the default password is subject to the eight-character limit. Any password set by the user can have a length of 8 to 40 characters.

```
login: admin
Password:
Please change your passwords now.
Use Control-C to exit or press 'Enter' key to proceed.
for user - root
Changing password for root
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - factory
Changing password for factory
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - admin
Changing password for admin
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
Please change your passwords now.
for user - user
Changing password for user
Enter new password: *****
Password changed.
Saving password to stable storage.
Password saved to stable storage successfully.
switch:admin>
```

Setting the date and time

Switches maintain the current date and time in flash memory. Date and time are used for logging events. Switch operation does not depend on the date and time; a switch with an incorrect date and time value still functions properly. Because the date and time are used for logging, set them correctly.

 **NOTE:** If secure mode is not enabled, a change in date or time to one switch is forwarded to the principal switch and distributed to the fabric. If secure mode is enabled, you can make date or time changes only on the primary FCS switch and then distribute the changes to the fabric.

Setting the date and time

1. Connect to the switch and log in as admin.
2. Issue the `date` command using the following syntax:

```
date "mmddHHMMyy"
```

where:

- *mm* is the month; valid values are 01 through 12.
- *dd* is the date; valid values are 01 through 31.
- *HH* is the hour; valid values are 00 through 23.
- *MM* is minutes; valid values are 00 through 59.
- *yy* is the year; valid values are 00 through 99 (values greater than 69 are interpreted as 1970–1999, and values less than 70 are interpreted as 2000–2069). For example:

```
switch:admin> date
Fri Jan 29 17:01:48 UTC 2000
switch:admin> date "0227123003"
Thu Feb 27 12:30:00 UTC 2003
switch:admin>
```

For details about changing time zones, see the `tsTimeZone` command in the *HP StorageWorks Fabric OS 5.x command reference guide*.

Synchronizing local time with an external source

To synchronize the local time of the principal or primary FCS switch to an external NTP server:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
tsclockserver "ipaddr"
```

where *ipaddr* is the IP address of the NTP server, which the switch must be able to access. This operand is optional; by default its value is `LOCL`, which uses the local clock of the principal or primary switch as the clock server. For example:

```
switch:admin> tsclockserver
LOCL
switch:admin> tsclockserver "132.163.135.131"
switch:admin> tsclockserver
132.163.135.131
switch:admin>
```

HP recommends that you synchronize time with an external NTP server, as described on ["Synchronizing local time with an external source"](#). If you cannot do so, use the next procedure.

Setting the time zone

1. Connect to the switch and log in as admin.
2. Issue the `tsTimeZone` command as follows:

```
tsTimeZone [houroffset [, minuteoffset]]
```

- For Pacific Standard Time enter `tsTimeZone -8,0`.
- For Central Standard Time enter `tsTimeZone -6,0`.
- For Eastern Standard Time enter `tsTimeZone -5,0`.

The default time zone for switches is Universal Time Conversion (UTC), which is 8 hours ahead of Pacific Standard Time. [Table 3](#) shows additional time zone conversion values.

The parameters do not apply if the time zone of the switch has already been changed from the default (8 hours ahead of Pacific Standard Time).

See the `tsTimeZone` command in the *HP StorageWorks Fabric OS 5.x command reference guide* for detailed information about the command parameters.

Repeat the procedure on all switches for which the Time Zone needs to be set. This needs to be done only once; the value is written to nonvolatile memory. For U.S. time zones, use [Table 3](#) to determine the correct parameter for the `tsTimeZone` command.

Table 3 Conversion from UTC to local time


Local time	Difference from UTC for <code>tsTimeZone</code>
Atlantic Standard	-4, 0
Atlantic Daylight	-3, 0
Eastern Standard	-5, 0
Eastern Daylight	-4, 0
Central Standard	-6, 0
Central Daylight	-5, 0
Mountain Standard	-7, 0
Mountain Daylight	-6, 0
Pacific Standard	-8, 0
Pacific Daylight	-7, 0
Alaskan Standard	-9, 0
Alaskan Daylight	-8, 0
Hawaiian Standard	-10, 0

Maintaining licensed features

Feature licenses might be part of the licensed Power Pack supplied with switch software, or you can purchase licenses separately from your switch vendor, who will provide you with keys to unlock the features. License keys are provided on a per-chassis basis, so for products that support multiple logical switches (domains), a license key applies to all domains within the chassis.

To unlock a licensed feature, you can either use the license key provided in the Power Pack or execute the following procedure to generate a license key at the HP web site:

<http://webkey.external.hp.com/welcome.asp>.

 **NOTE:** For each chassis to be licensed, you need a transaction key and a license ID. The transaction key is in the Power Pack supplied with the switch software; or, when you purchase a license, your switch vendor gives you a transaction key to be used to obtain a license key. To see a switch license ID, issue the `licenseIdShow` command.

Unlocking a licensed feature

1. If you already have a license key, go to [step 10](#).
If you do not have a key, launch an Internet browser and visit the HP web site:
<http://webkey.external.hp.com/welcome.asp>.
2. Click **products**.
3. Click **Software Products**.
4. In the Related Links panel on the right side of the page, select **Software License Keys**.
The Software License Keys instruction page appears.

5. If you want to generate a single license key, select **Generate 1 license key**.
If you want to generate multiple license keys, select **Batch Generation of Licenses**.
The Software License Key instruction page opens.
6. Enter the requested information in the required fields.
When generating multiple license keys, enter the worldwide names and transaction keys in the table at the bottom of the screen. If you need additional rows in the table, select **Add More Rows**.
7. Click **Next**.
A verification screen appears.
8. Verify that the information is correct.
Click **Submit** if the information displayed is correct. If the information is incorrect, click **Previous** and change the information.
9. After the information is corrected, click **Submit**.
An information screen displays the license keys. You also receive an e-mail from the HP licensing company.
10. Activate and verify the license as follows:
 - a. Connect to the switch and log in as admin.
 - b. Activate the license using the `licenseAdd` command:

```
switch:admin> licenseadd "key"
```

The license key is case-sensitive and must be entered exactly as given. The quotation marks are optional.

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, licenses are effective on both CP blades and on all logical switches, but are valid only when the CP blade is inserted into a chassis that has an appropriate license ID stored in the World Wide Name (WWN) card. If a CP is moved from one chassis to another, the license works in the new chassis only if the WWN card is the same in the new chassis. Otherwise, a new license key is generated.

For example, if you swap one CP blade at a time, or replace a single CP blade, the existing CP blade (the active CP blade) propagates the licenses to the new CP blade.

If you move a standby CP from one chassis to another, the active CP propagates its configuration (including license keys).
 - c. Verify that the license was added by issuing the `licenseShow` command.
The licensed features currently installed on the switch are listed. If the feature is not listed, reissue the `licenseAdd` command.
 - d. Some features may require additional configuration, or you might need to disable and reenable the switch to make them operational; see the feature documentation for details. For example:

```
switch:admin> licenseshow
SbeSdQdQySyrTeJ:
Web license
Zoning license
Fabric license
Remote Switch license
Extended Fabric license
Fabric Watch license
Performance Monitor license
Trunking license
Security license
SbbebdQS9QTscfcB:
Ports on Demand license - additional 8 port upgrade
SbbebdQS9QTcgfcz:
Ports on Demand license - additional 8 port upgrade
```

Removing a licensed feature

1. Connect to the switch and log in as admin.
2. Issue the `licenseShow` command to display the active licenses.
3. Remove the license key using the `licenseRemove` command:

```
switch:admin> licenseremove "key"
```

The license key is case-sensitive and must be entered exactly as given. The quotation marks are optional. After removing a license key, the optionally licensed feature is disabled when the switch is rebooted or when a `switchDisable` or `switchEnable` is performed.

4. Issue the `licenseShow` command to verify that the license is disabled. For example:

```
switch:admin> licenseshow
bQebzbRdScRfc0iK:
  Web license
  Zoning license
SybbzQQ9edTzcc0X:
  Fabric license
switch:admin> licenseremove "bQebzbRdScRfc0iK"
removing license key "bQebzbRdScRfc0iK"
switch:admin>
```

After a reboot (or `switchDisable` and `switchEnable`):

```
switch:admin> licenseshow
SybbzQQ9edTzcc0X:
  Fabric license
switch:admin>
```

If there are no license keys, `licenseShow` displays the message `No licenses`.


Customizing the switch name

Switches can be identified by IP address, Domain ID, WWN, or by customized switch names that are unique and meaningful.

Version 4.0.0 (and later) switch names can be from 1 to 15 characters long, must begin with a letter, and can contain letters, numbers, or the underscore character. It is not necessary to use quotation marks.

The default names are the following:

- For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32: the name is `swd77`.
- For the Core Switch 2/64: the name varies depending on the number of logical switches. The two logical switches have different default names. The name `swd77` is used for the logical switch containing the port blades in slots 1 through 4; `swd76` is used for the logical switch containing the port blades in slots 7 through 10.
- For the SAN Director 2/128 and the 4/256 SAN Director: the name is `swd77`.

 **NOTE:** Changing the switch name causes a domain address format Registered State Change Notification (RSCN) to be issued.

Customizing the switch name

1. For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32: Proceed to the next step.

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director: Open a telnet window for each logical switch and issue the `switchName` command.

2. Connect to the switch and log in as admin.
3. For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32: Proceed to the next step.
For the SAN Director 2/128 and 4/256 SAN Director: If configured for one domain (the default) proceed to the next step. If configured with two domains, proceed as for the Core Switch 2/64.
For the Core Switch 2/64: Choose the logical switch that you want to change. Enter the value that corresponds to that logical region:
 - Enter 0 to configure logical switch 0 (slot 1 through 4).
 - Enter 1 to configure logical switch 1 (slot 7 through 10).
4. Issue the `switchName` command with the following syntax:

```
switchname "newname"
```

 where *newname* is the new name for the switch.
5. Record the new switch name for future reference.
6. For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director configured with two domains: Disconnect from the session and repeat the procedure for the second logical switch. For example:

```
switch:admin> switchname "switch62"
Committing configuration...
Done.
switch62:admin>
```

Customizing the chassis name

Beginning with Fabric OS 4.4.0, HP recommends that you customize the chassis name for each switch. Some system logs identify switches by chassis names, so if you assign meaningful chassis names in addition to meaningful switch names, logs are more useful.

Changing the chassis name

1. Connect to the switch and log in as admin.
2. Issue the `chassisName` command, with the following syntax:

```
chassisname "newname"
```

 where *newname* is the new name for the chassis.
 Chassis names can contain 1 to 15 characters, must begin with a letter, and can consist of letters, numerals, and the underscore character. The quotation marks are optional.
3. Record the new chassis name for future reference.

Disabling and enabling a switch

By default, the switch is enabled after power is applied and diagnostics and switch initialization routines have finished. You can disable and reenable it as necessary.

Disabling a switch

1. Connect to the switch and log in as admin.
2. Issue the `switchDisable` command.
 All Fibre Channel ports on the switch are taken offline. If the switch was part of a fabric, the fabric reconfigures.

Enabling a switch

1. Connect to the switch and log in as admin.
2. Issue the `switchEnable` command.
 All Fibre Channel ports that pass the Power-on Self Test (POST) are enabled. If the switch has interswitch links (ISLs) to a fabric, it joins the fabric.

Disabling and enabling a port

All licensed ports are enabled by default. You can disable and reenable them as necessary. Ports that you activate with Ports on Demand must be enabled explicitly, as described in ["Activating Ports on Demand"](#) on page 30.

Disabling a port

1. Connect to the switch and log in as admin.
2. For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32: Issue the following command:

```
portdisable portnumber
```

where *portnumber* is the port number of the port you want to disable.

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director: Issue the following command:

```
portdisable slotnumber/portnumber
```

where *slotnumber* and *portnumber* are the slot and port numbers of the port you want to disable.

 **NOTE:** If the port is connected to another switch, the fabric might reconfigure.

Enabling a port

1. Connect to the switch and log in as admin.
2. For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32: Issue the following command:


```
portenable portnumber
```

where *portnumber* is the port number of the port you want to enable.

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director: Issue the following command:

```
portenable slotnumber/portnumber
```

where *slotnumber* and *portnumber* are the slot and port numbers of the port you want to enable. (Slots are numbered 1 through 4 and 7 through 10, counting from left to right.)

 **NOTE:** If the port is connected to another switch, the fabric might reconfigure. If the port is connected to one or more devices, these devices become available to the fabric.

If you change port configurations during a switch failover, the ports might become disabled. To bring the ports online, reissue the `portEnable` command after the failover is complete.

Activating Ports on Demand

The SAN Switch 4/32 can be purchased with 16, 24, or 32 licensed ports. As your needs increase, you can activate unlicensed ports (up to the maximum of 32 ports) by purchasing and installing the HP Ports on Demand optional, licensed product.

The 4/8 SAN Switch and 4/16 SAN Switch can be purchased with 8 ports and no E_Port, 8 ports with full-fabric access (4/8 SAN Switch), and 16 ports with full-fabric access (4/16 SAN Switch). If you purchase the 4/8 SAN Switch with 8 ports enabled, you can activate unlicensed ports in 4-port increments up to 16 ports by purchasing and installing the HP StorageWorks 4/8 SAN 4-Port Upgrade License. You can also purchase a full-fabric upgrade license if your switch does not support full-fabric access.

Ports on Demand is ready to be unlocked in the switch firmware. Its license key might be part of the licensed software supplied with your switch, or you can purchase the license key separately from your

switch vendor. You might need to generate a license key from a transaction key supplied with your purchase. If so, launch an Internet browser and visit the HP web site: <http://webkey.external.hp.com/welcome.asp>. Select **Generate a license key** and follow the instructions to generate the key.

By default, ports 0 through 15 are activated on the SAN Switch 4/32. Each Port upgrade license activates the next group of eight ports in numerical order. Before installing a license key, you must insert transceivers in the ports to be activated. Remember to insert the transceivers in the lowest group of inactive port numbers first. For example, if only 16 ports are currently active and you are installing one 8-Port Upgrade License key, make sure to insert the transceivers in ports 16 through 23. If you later install a second license key, insert the transceivers in ports 24 through 31. For details on inserting transceivers, see the *HP StorageWorks SAN Switch 4/32 installation guide*.

After you install a license key, you must enable the ports to complete their activation. You can do so without disrupting switch operation by issuing the `portEnable` command on each port. You can also disable and reenable the switch to activate ports.

Activating Ports on Demand

1. Connect to the switch and log in as admin.
2. Optional: To verify the current states of the ports, issue the `portShow` command.
In the `portShow` output, the Licensed field indicates whether or not the port is licensed.
3. Install the HP Port Upgrade License.
For instructions, see "[Maintaining licensed features](#)" on page 26.
4. Issue the `portEnable` command to enable the ports.
5. Optional: issue the `portShow` command to check the newly activated ports.

If you remove a Port Upgrade License, the licensed ports become disabled after the next platform reboot or the next port deactivation.

Making basic connections

You can make basic connections to devices and to other switches.

Before connecting a version 4.0.0 (or later) switch to a fabric that contains switches running earlier firmware versions, you must first set the same port identifier (PID) format on all the switches. The presence of different PID formats in a fabric causes fabric segmentation.

For information on PID formats and related procedures, see "[Selecting a PID format](#)" on page 215.

For information on configuring the routing of connections, see "[Routing traffic](#)" on page 97.

For information on configuring extended interswitch connections, see "[Administering extended fabrics](#)" on page 163.

Connecting to devices

To minimize port logins, power off all devices before connecting them to the switch. For devices that cannot be powered off, first use the `portDisable` command to disable the port on the switch, and then connect the device. When powering the devices back on, wait for each device to complete the fabric login before powering on the next one.

Connecting to other switches

See the SAN Switch installation guide for your switch model for ISL connection and cable management information. The standard (default) ISL mode is L0, which you can configure with the `portCfgLongDistance` command. ISL Mode L0 is a static mode, with the following maximum ISL distances:

- 10 km at 1 Gbit/second
- 5 km at 2 Gbit/second
- 2.5 km at 4 Gbit/second

ISL mode L0 is available on all Fabric OS releases. When you upgrade from Fabric OS 4.0.0 to Fabric 4.1.0 or later, all extended ISL ports are set to L0 mode.

For information on extended ISL modes, which enable longer-distance ISLs, see ["Administering extended fabrics"](#) on page 163.

Working with domain IDs

Although domain IDs are assigned dynamically when a switch is enabled, you can reset them manually to control the ID number or to resolve a domain ID conflict when you merge fabrics.

If a switch already has a domain ID when it is enabled, and that domain ID conflicts with a switch already in the fabric, the conflict is resolved. The process can take several seconds, during which traffic is delayed.

The default domain ID for HP StorageWorks switches is 1. The default domain ID applies to the logical switches in the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director that are configured for two domains. To prevent domain conflict, you can either disable one of the switches until the other is connected to the fabric, and then reenabling the switches so that unique domain IDs are assigned, or you can use the procedure ["Setting the domain ID"](#) on page 33 to make the domain IDs unique before connecting the logical switches to the fabric.

△ **CAUTION:** On switches running Fabric OS 4.0.0 and later, do not use domain ID 0, which is reserved for another purpose. The use of this domain ID can cause the switch to reboot continuously.

Avoid changing the domain ID on the FCS in secure mode. To minimize down time, change the domain IDs on the other switches in the secure fabric.

Displaying domain IDs

1. Connect to a switch and log in as admin.
2. Issue the `fabricShow` command.

Fabric information is displayed, including the domain ID (D_ID), for example:

```
ras045:admin> fabricshow
Switch ID      Worldwide Name      Enet IP Addr      FC IP Addr      Name
-----
1: fffc01 10:00:00:60:69:e4:00:3c 10.32.220.80      0.0.0.0      "ras080"
2: fffc02 10:00:00:60:69:e0:01:46 10.32.220.1      0.0.0.0      "ras001"
3: fffc03 10:00:00:60:69:e0:01:47 10.32.220.2      0.0.0.0      "ras002"
5: fffc05 10:00:00:05:1e:34:01:bd 10.32.220.5      0.0.0.0      "ras005"
6: fffc06 10:00:00:05:1e:34:02:3e 10.32.220.6      0.0.0.0      "ras006"
7: fffc07 10:00:00:60:69:34:02:0c 10.32.220.7      0.0.0.0      "ras007"
10: fffc0a 10:00:00:60:69:80:04:46 10.32.220.10     10.32.219.0    "ras010"
11: fffc0b 10:00:00:60:69:80:04:47 10.32.220.11     10.32.219.1    "ras011"
15: fffc0f 10:00:00:60:69:80:47:74 10.32.220.15     0.0.0.0      "ras015"
16: fffc10 10:00:00:60:69:80:47:75 10.32.220.16     0.0.0.0      "ras016"
19: fffc13 10:00:00:05:1e:34:00:ad 10.32.220.19     0.0.0.0      "ras019"
20: fffc14 10:00:00:05:1e:34:00:63 10.32.220.20     0.0.0.0      >"ras020"
30: fffc1e 10:00:00:60:69:90:02:21 10.32.220.30     0.0.0.0      "ras030"
31: fffc1f 10:00:00:60:69:90:02:60 10.32.220.31     0.0.0.0      "ras031"
32: fffc20 10:00:00:60:69:90:02:68 10.32.220.32     0.0.0.0      "ras032"
33: fffc21 10:00:00:60:69:90:03:20 10.32.220.33     0.0.0.0      "ras033"
34: fffc22 10:00:00:60:69:90:03:01 10.32.220.34     0.0.0.0      "ras034"
40: fffc28 10:00:00:60:69:50:06:7f 10.32.220.40     0.0.0.0      "ras040"
45: fffc2d 10:00:00:05:1e:34:c5:17 10.32.220.45     0.0.0.0      "ras045"
50: fffc32 10:00:00:60:69:c0:06:64 10.32.220.50     0.0.0.0      "ras050"
51: fffc33 10:00:00:60:69:c0:1e:43 10.32.220.51     0.0.0.0      "ras051"
60: fffc3c 10:00:00:60:69:12:34:44 10.32.220.60     0.0.0.0      "ras060"
62: fffc3e 10:00:00:60:69:12:32:76 10.32.220.62     0.0.0.0      "ras062"
63: fffc3f 10:00:00:60:69:12:45:6e 10.32.220.63     0.0.0.0      "ras063"
64: fffc40 10:00:00:60:69:12:1d:51 10.32.220.64     0.0.0.0      "ras064"

The Fabric has 25 switches

ras045:admin>
```

The fields in the `fabricShow` display are:

- Switch ID: The switch Domain_ID and embedded port D_Id.
- Worldwide Name: The switch WWN.
- Enet IP Addr: The switch Ethernet IP address.
- FC IP Addr: The switch FC IP address.
- Name: The switch symbolic name. An arrow (>) indicates the principal switch.

Setting the domain ID

1. Connect to the switch and log in as admin.
2. Issue the `switchDisable` command to disable the switch.
3. Issue the `configure` command.
4. Enter `y` after the Fabric Parameters prompt:

```
Fabric parameters (yes, y, no, n): [no] y
```
5. Enter a unique domain ID at the Domain prompt. Use a domain ID value from 1 through 239 for normal operating mode (FCSW compatible):

```
Domain: (1..239) [1] 3
```
6. Respond to the remaining prompts (or press **Ctrl-d** to accept the other settings and exit).
7. Issue the `switchEnable` command to reenabale the switch.

Linking through a gateway

A gateway merges SANs into a single fabric by establishing point-to-point E_Port connectivity between two Fibre Channel switches that are separated by a network with a protocol such as IP or SONET.

Except for link initialization, gateways are transparent to switches; the gateway simply provides E_Port connectivity from one switch to another.

By default, switch ports initialize links using the Exchange Link Parameters (ELP) mode 1. However, gateways expect initialization with ELP mode 2 (also called *ISL R_RDY mode*). Therefore, to enable two switches to link through a gateway, the ports on both switches must be set for ELP mode 2.

Any number of E_Ports in a fabric can be configured for gateway links, provided the following rules are followed:

- All switches in the fabric must be upgraded to Fabric OS 3.1.0 or later, or to 4.1.0 or later.
- To prevent fabric segmentation, make sure that all switches in the fabric are using the core PID format, as described in section “[Configuring a link through a gateway](#)” next.
- When determining switch count maxima, include the switches connected to both sides of the gateway.
- Extended links (those created using the Extended Fabrics licensed feature) and the security features in Secure Fabric OS are not supported through gateway links.

Configuring a link through a gateway

1. If you are not sure that the PID format is consistent across the entire fabric, issue the `configShow` command on all switches to check the PID settings. If necessary, change the PID format on any non-conforming switches, as described in “[Configuring the PID format](#)” on page 213.
2. Connect to the switch on one end of the gateway and log in as admin.

3. Issue the `portCfgIslMode` command:

For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32:

```
portCfgIslMode port mode
```

Specify a port number. Valid values for port number depend on the switch type. The mode operand is required: Specify 1 to enable ISL R_RDY mode (gateway link) or specify 0 to disable it.

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director:

```
portCfgIslMode slot/port, mode
```

Specify a slot/port number pair. Valid values for slot and port number depend on the switch type. The mode operand is required: Specify 1 to enable ISL R_RDY mode (gateway link) or specify 0 to disable it.

In the following example, slot 2, port 3 is enabled for a gateway link:

```
switch:admin> portcfgislmode 2/3, 1
Committing configuration...done.
ISL R_RDY Mode is enabled for port 3. Please make sure the PID
formats are consistent across the entire fabric.
switch:admin>
```

4. Repeat the steps for any additional ports to be connected to the gateway.
5. Repeat the procedure on the switch at the other end of the gateway.

See the *HP StorageWorks Fabric OS 5.x command reference guide* for more information about the `portCfgIslMode` command.

Checking status

You can check the status of switch operation, high availability (HA) features, and fabric connectivity.

Checking switch operation

1. Connect to the switch and log in as admin.
2. Issue the `switchShow` command.
A switch summary and a port summary are displayed.
3. Verify that the switch and ports are online.
4. Issue the `switchStatusShow` command to further check the status of the switch.

Verifying HA features

HA features provide maximum reliability and non-disruptive replacement of key hardware and software modules. To verify these features, connect to the switch as admin and use any of the following commands:

- The `chassisshow` verifies the field replaceable units (FRUs).
- For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director:
 - The `haShow` command verifies that HA is enabled, that the heartbeat is up, and that the HA state is synchronized between the active and standby CP blades.
 - The `slotShow` command inventories and displays the current status of each slot in the system.

Verifying fabric connectivity

1. Connect to the switch and log in as admin.
2. Issue the `fabricShow` command.
A summary of all the switches in the fabric is displayed. For example:

```
switch:admin> fabricshow
```

Switch ID	Worldwide Name	Enet IP Addr	FC IP Addr	Name
1: fffc01	10:00:00:60:69:80:04:5a	192.168.186.61	192.168.68.193	"switch61"
3: fffc03	10:00:00:60:69:10:9c:29	192.168.186.175	0.0.0.0	"switch175"
4: fffc04	10:00:00:60:69:12:14:b7	192.168.174.70	0.0.0.0	"switch70"
5: fffc05	10:00:00:60:69:45:68:04	192.168.144.121	0.0.0.0	"switch121"
6: fffc06	10:00:00:60:69:00:54:ea	192.168.174.79	192.168.68.197	"switch79"
7: fffc07	10:00:00:60:69:80:04:5b	192.168.186.62	192.168.68.194	"switch62"
8: fffc08	10:00:00:60:69:04:11:22	192.168.186.195	0.0.0.0	"switch195"
9: fffc09	10:00:00:60:69:10:92:04	192.168.189.197	192.168.68.198	"switch197"
10: fffc0a	10:00:00:60:69:50:05:47	192.168.189.181	192.168.68.181	"switch181"
11: fffc0b	10:00:00:60:69:00:54:e9	192.168.174.78	192.168.68.196	"switch78"
15: fffc0f	10:00:00:60:69:30:1e:16	192.168.174.73	0.0.0.0	"switch73"
33: fffc21	10:00:00:60:69:90:02:5e	192.168.144.120	0.0.0.0	"switch120"
44: fffc2c	10:00:00:60:69:c0:06:8d	192.168.144.121	0.0.0.0	"switch121"
97: fffc61	10:00:00:60:69:90:02:ed	192.168.144.123	0.0.0.0	"switch123"
98: fffc62	10:00:00:60:69:90:03:32	192.168.144.122	0.0.0.0	"switch122"

```

The Fabric has 15 switches

switch:admin>

```

Verifying device connectivity

1. Connect to the switch and log in as admin.
2. Optional: Issue the `switchShow` command to verify that devices, hosts, and storage are connected.
3. Optional: Issue the `nsShow` command to verify that devices, hosts, and storage have successfully registered with the Name Server.
4. Issue the `nsAllShow` command.

The 24-bit Fibre Channel addresses of all devices in the fabric are displayed. For example:

```
switch:admin> nsallshow
{
  010e00 012fe8 012fef 030500 030b04 030b08 030b17 030b18
  030b1e 030b1f 040000 050000 050200 050700 050800 050de8
  050def 051700 061c00 071a00 073c00 090d00 0a0200 0a07ca
  0a07cb 0a07cc 0a07cd 0a07ce 0a07d1 0a07d2 0a07d3 0a07d4
  0a07d5 0a07d6 0a07d9 0a07da 0a07dc 0a07e0 0a07e1 0a0f01
  0a0f02 0a0f0f 0a0f10 0a0f1b 0a0f1d 0b2700 0b2e00 0b2fe8
  0b2fef 0f0000 0f0226 0f0233 0f02e4 0f02e8 0f02ef 210e00
  211700 211fe8 211fef 2c0000 2c0300 611000 6114e8 6114ef
  611600 620800 621026 621036 6210e4 6210e8 6210ef 621400
  621500 621700 621a00
  75 Nx_Ports in the Fabric }
switch:admin>
```

The number of devices listed should agree with the number of devices that are connected.

Tracking and controlling switch changes

The TC feature allows you to keep a record of specific changes that might not be considered switch events, but might provide useful information. The output from the TC feature is dumped to the system messages log for the switch. Use the `errDump` or `errShow` command to view the log.

Items in the log created from the TC feature are labeled `TRCK`.

Trackable changes are:

- Successful login
- Unsuccessful login
- Logout

- Configuration file change from task
- TC feature on
- TC feature off

An SNMP-TRAP mode can also be enabled; see the `trackChangesHelp` command in the *HP StorageWorks Fabric OS 5.x command reference guide*.

For troubleshooting information on the TC feature, see ["Inaccurate information in the system message log"](#) on page 161.

Enabling the TC feature

1. Connect to the switch and log in as admin.
2. Issue the `trackChangesSet 1` command to enable the TC feature:

A message is displayed, verifying that the TC feature is on:

```
switch:admin> trackchangesset 1
Committing configuration...done.
switch:admin>
```

The output from the TC feature is dumped to the system message log for the switch. Use the `errDump` or `errShow` command to view the log.

Items in the system message log created from the TC feature are labeled TRCK:

```
2004/08/24-08:45:43, [TRCK-1001], 212,, INFO, ras007, Successful login by user
admin.
```

Displaying the status of the TC feature

1. Connect to the switch and log in as admin.
2. Issue the `trackChangesShow` command.

The status of the TC feature is displayed as either on or off. The display tells whether the TC feature is configured to send SNMP traps:

```
switch:admin> trackchangesshow
Track changes status: ON
Track changes generate SNMP-TRAP: NO
switch:admin>
```

Viewing the switch status policy threshold values

1. Connect to the switch and log in as admin.
2. Issue the `switchStatusPolicyShow` command.

Whenever there is a switch change, an error message is logged and an SNMP `connUnitStatusChange` trap is sent.

For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32: The output is similar to the following:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
                Down      Marginal
-----
PowerSupplies  2          1
Temperatures   2          1
Fans           2          1
Flash          0          1
MarginalPorts  5          2
FaultyPorts    2          1
MissingSFPs    2          1
switch:admin>
```

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director: The output is similar to the following:

```
switch:admin> switchstatuspolicyshow
The current overall switch status policy parameters:
              Down      Marginal
-----
PowerSupplies 3         0
Temperatures  2         1
Fans          2         1
WWN           0         1
CP            0         1
Blade         0         1
Flash         0         1
MarginalPorts 2         1
FaultyPorts   2         1
MissingSFPs   0         0
switch:admin>
```


The policy parameter determines the number of failed or inoperable units for each contributor that triggers a status change in the switch.

Each parameter can be adjusted so that a specific threshold must be reached before that parameter changes the overall status of a switch to MARGINAL or DOWN. For example, if the `FaultyPorts` DOWN parameter is set to 3, the status of the switch changes if 3 ports fail. Only one policy parameter needs to pass the MARGINAL or DOWN threshold to change the overall status of the switch.

For more information about setting policy parameters, see the *HP StorageWorks Fabric OS 5.x administrator guide*.

Setting the switch status policy threshold values

1. Connect to the switch and log in as admin.
2. Issue the `switchStatusPolicySet` command.
The current switch status policy parameter values are displayed first. You are then prompted to enter values for each DOWN and MARGINAL threshold parameter:
3. Verify the threshold settings you have configured for each parameter by issuing the `switchStatusPolicyShow` command to view your current switch status policy configuration.

 **NOTE:** By setting the DOWN and MARGINAL values for a parameter to 0, 0 that parameter is no longer used in setting the overall status for the switch.

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32: The following example shows the command as executed on a SAN Switch 2/32. The output is similar on SAN Switch 2/8V, SAN Switch 2/16V, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32:

```
switch:admin> switchstatuspolicyset
To change the overall switch status policy parameters
The current overall switch status policy parameters:
      Down      Marginal
-----
FaultyPorts  2          1
MissingSFPs  0          0
PowerSupplies 2          1
Temperatures 2          1
Fans         2          1
PortStatus   0          0
ISLStatus    0          0
Note that the value, 0, for a parameter, means that it is
NOT used in the calculation.
** In addition, if the range of settable values in the prompt is (0..0),
** the policy parameter is NOT applicable to the switch.
** Simply hit the Return key.
The minimum number of
FaultyPorts contributing to
                        DOWN status: (0..32) [2] 3
FaultyPorts contributing to
                        MARGINAL status: (0..32) [1] 2
MissingSFPs contributing to
                        DOWN status: (0..32) [0]
MissingSFPs contributing to
                        MARGINAL status: (0..32) [0]
Bad PowerSupplies contributing to
                        DOWN status: (0..2) [2]
Bad PowerSupplies contributing to
                        MARGINAL status: (0..2) [1]
Bad Temperatures contributing to
                        DOWN status: (0..5) [2]
Bad Temperatures contributing to
                        MARGINAL status: (0..5) [1]
Bad Fans contributing to
                        DOWN status: (0..6) [2]
Bad Fans contributing to
                        MARGINAL status: (0..6) [1]
Down PortStatus contributing to
                        DOWN status: (0..32) [0]
Down PortStatus contributing to
                        MARGINAL status: (0..32) [0]
down ISLStatus contributing to
                        DOWN status: (0..32) [0]
down ISLStatus contributing to
                        MARGINAL status: (0..32) [0]
Policy parameter set has been changed
```

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director: Command output includes parameters related to CP blades.

3 Configuring standard security features

This chapter provides information and procedures for configuring standard Fabric OS security features such as account and password management.

Additional security features are available when secure mode is enabled. For information about licensed security features available in Secure Fabric OS, see the *HP StorageWorks Secure Fabric OS administrator guide*.

Secure protocols

Fabric OS supports the secure protocols shown in [Table 4](#).

Table 4 Secure protocol support

Protocol	Description
Secure Sockets Layer (SSL)	Supports SSLv3, 128-bit encryption by default. Fabric OS uses SSL to support HTTPS. A certificate must be generated and installed on each switch to enable SSL.
HTTPS	Advanced Web Tools supports the use of HTTPS.
Secure File Copy (scp)	Configuration upload and download support the use of scp.
SNMPv3	SNMPv1 is also supported.

SNMP is a standard method for monitoring and managing network devices. Using SNMP components, you can program tools to view, browse, and manipulate HP StorageWorks switch variables and set up enterprise-level management processes.

Every HP StorageWorks switch carries an SNMP agent and Management Information Base (MIB). The agent accesses MIB information about a device and makes it available to a network manager station. You can manipulate information of your choice by trapping MIB elements using the Fabric OS CLI, Advanced Web Tools, or Fabric Manager.

The SNMP Access Control List (ACL) provides a way for you to restrict SNMP get and set operations to certain hosts and IP addresses. This is used for enhanced management security in the SAN.

For details on HP StorageWorks MIB files, naming conventions, loading instructions, and information about using the HP SNMP agent, see the *HP StorageWorks Fabric OS 5.x MIB reference guide*.

[Table 5](#) describes additional software or certificates that you must obtain to deploy secure protocols.

Table 5 Items needed to deploy secure protocols

Protocol	Host side	Switch side
Secure telnet (sectelnet)	Sectelnet client	License not required, but a switch certificate issued by HP is required
Secure Shell (SSH)	SSH client	None
HTTPS	No requirement on host side except a browser that supports HTTPS	Switch IP certificate for SSL
Secure File Copy (scp)	SSH daemon, scp server	None
SNMPv3, SNMPv1	None	None


The security protocols are designed with the four main usage cases described in [Table 6](#).

Table 6 Main security scenarios

Fabric	Management interfaces	Comments
Nonsecure	Nonsecure	No special setup is need to use telnet or HTTP. An HP switch certificate must be installed if sectelnet is used.
Nonsecure	Secure	Secure protocols may be used. An SSL switch certificate must be installed if SSH/HTTPS is used.
Secure	Secure	<p>Secure protocols are supported on Fabric OS 4.4.0 (and later) switches. Switches running earlier Fabric OS versions can be part of the secure fabric, but they do not support secure management.</p> <p>Secure management protocols must be configured for each participating switch. Nonsecure protocols may be disabled on nonparticipating switches.</p> <p>If SSL is used, certificates must be installed.</p>
Secure	Nonsecure	<p>You must use sectelnet because telnet is not allowed in secure mode.</p> <p>Nonsecure management protocols are necessary under these circumstances:</p> <ul style="list-style-type: none">• The fabric contains switches running Fabric OS 3.2.0.• The presence of software tools that do not support Secure protocols: for example, Fabric Manager 4.0.0.• The fabric contains switches running Fabric OS versions earlier than 4.4.0. Nonsecure management is enabled by default.

Ensuring network security

To ensure security, Fabric OS supports SSH encrypted sessions. SSH encrypts all messages, including the client's transmission of password during login. The SSH package contains a daemon (sshd), which runs on the switch. The daemon supports a wide variety of encryption algorithms, such as Blowfish-CBC and AES.

 **NOTE:** To maintain a secure network, avoid using telnet or any other unprotected application when you are working on the switch. For example, if you use telnet to connect to a machine, and then start an SSH or secure telnet session from that machine to the switch, the communication to the switch is in clear text and, therefore, is not secure.

Nor is the FTP protocol secure. When you use FTP to copy files to or from the switch, the contents are in clear text. When you use FTP to copy files to or from the switch, the contents, including the remote FTP server's login and password, are in clear text. This limitation affects the following commands: `saveCore`, `configUpload`, `configDownload`, and `firmwareDownload`.


Commands that require a secure login channel must be issued from an original SSH session. If you start an SSH session, and then use the `login` command to start a nested SSH session, commands that require a secure channel are rejected.

Fabric OS 4.4.0 and later supports SSH protocol 2.0 (ssh2). For more information on SSH, see the SSH IETF web site: <http://www.ietf.org/ids.by.wg/secsh.html>. Another informative source is *SSH, The Secure Shell: The Definitive Guide* by Daniel J. Barrett, Richard Silverman.

Fabric OS 4.4.0 comes with the SSH server preinstalled; however, you must select and install the SSH client. For information on installing and configuring the F-Secure SSH client, visit the following web site: <http://www.f-secure.com>.

Configuring the telnet interface

Telnet is enabled by default. To prevent users from passing clear text passwords over the network when they connect to the switch, you can disable the telnet interface.

 **NOTE:** Before disabling the telnet interface, make sure you have an alternate method of establishing a connection with the switch.

Disabling telnet

1. Connect to the switch and log in as admin.

Connect through some other means than telnet, for example, through SSH.

2. Issue the following command:

```
configure telnetd
```

3. In response to the System Services prompt, enter `y`.

4. In response to the telnetd prompt, enter `off`.

The telnet interface is disabled. If you entered the command during a standard telnet session, the session terminates. For example:

```
switch:admin> configure telnetd
Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.
Configure...
ssl attributes (yes, y, no, n): [no]
http attributes (yes, y, no, n): [no]
snmp attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
cfgload attributes (yes, y, no, n): [no]

[31454]: Read 1 license entries for generation 1.
[31454]: Read 1 license records.
System services (yes, y, no, n): [no] y

rstatd (on, off): [off]
rusersd (on, off): [off]
telnetd (on, off): [on] off
```

Enabling telnet

1. Connect to the switch through a means other than telnet (for example, SSH) and log in as admin.

2. Issue the following command:

```
configure telnetd
```

3. In response to the System Services prompt, enter `y`.

4. In response to the `telnetd` prompt, enter `on`.
The telnet interface is enabled.

Blocking listeners

HP StorageWorks switches block Linux® subsystem listener applications that are not used to implement supported features and capabilities. Table 7 lists the listener applications that HP StorageWorks switches either block or do not start.

Table 7 Blocked listener applications

Listener application	Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director	4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/8-EL, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32
chargen	Do not start	Do not start
echo	Do not start	Do not start
daytime	Do not start	Do not start
discard	Do not start	Do not start
ftp	Do not start	Do not start
rexec	Block with packet filter	Do not start
rsh	Block with packet filter	Do not start
rlogin	Block with packet filter	Do not start
time	Block with packet filter	Do not start
rstats	Do not start	Do not start
rusers	Do not start	Do not start

Accessing switches and fabrics

Table 8 lists the defaults for accessing hosts, devices, switches, and zones.

Table 8 Access defaults

Category	Default
Hosts	Any host can access the fabric by SNMP. Any host can telnet to any switch in the fabric. Any host can establish an HTTP connection to any switch in the fabric. Any host can establish an API connection to any switch in the fabric.
Devices	All device ports can access SES. All devices can access the management server. Any device can connect to any FC port in the fabric.
Switch access	Any switch can join the fabric. All switches in the fabric can be accessed through serial port.
Zoning	Node WWNs can be used for WWN-based zoning.

Creating and maintaining user-defined accounts

In addition to the default administrative and user accounts, Fabric OS supports up to 15 user-defined accounts in each logical switch (domain). These accounts expand your ability to track account access and audit administrative activities.


User-defined accounts can be assigned either admin-, switchAdmin-, or user-level roles. Admin-level accounts allow up to two simultaneous login sessions. User-level accounts allow up to four simultaneous login sessions. The total number of simultaneous login sessions allowed per logical switch is 15.

You can change passwords on user-defined accounts as described in ["Changing an account password"](#) on page 45.

If the TC feature is enabled, the system keeps track of account names and login attempts. (See ["Tracking and controlling switch changes"](#) on page 35 for details on enabling the TC feature.)

For large enterprises, Fabric OS also supports RADIUS services, as described in ["Setting up RADIUS AAA service"](#) on page 45.

The following procedures are for operations you can perform on user-defined accounts.

 **NOTE:** If you are operating in secure mode, you can perform these operations only on the primary FCS switch.

Displaying account information

1. Connect to the switch and log in as admin.
2. Issue one of the following commands:
 - `userConfig --show -a` to show all account information for a logical switch
 - `userConfig --show -b` to show all backup account information for a logical switch
 - `userConfig --show username` to show account information for the specified account name

Accounts with the admin role can display information about all accounts on the logical switch. Accounts with the switchAdmin role can display information about all accounts on the logical switch; however, they cannot display information about security, user management, or zoning. Accounts with the user role can display information only about themselves.

Creating a user-defined account

Accounts with the admin role can create accounts. Accounts with the user role cannot.

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
userConfig --add username -r rolename [-d description]
```

where:

<code>username</code>	Specifies the account name, which must begin with an alphabetic character. The name can consist of 8 to 40 characters. It is case-sensitive and can contain alphabetic and numeric characters, the dot, and the underscore. It must be different from all other account names on the logical switch.
<code>-r rolename</code>	Specifies the role: either admin, switchAdmin, or user in nonsecure mode; admin, user, or nonfcsadmin in secure mode.
<code>-d description</code>	Is an optional argument that adds a description to the account. The description field can be up to 40 printable ASCII characters. The following characters are not allowed: asterisk (*), quotation mark ("), exclamation point (!), semicolon (;), and colon (:).

3. In response to the prompt, enter a password for the account.
The password is not displayed when you enter it on the command line.

Deleting a user-defined account

Only accounts with the admin role can delete user-defined accounts on the logical switch.

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
userConfig --delete username
```

where *username* specifies the account name. You cannot delete the default accounts. An account cannot delete itself. All active CLI sessions for the deleted account are logged out.

3. Enter *y* at the prompt for confirmation.

Changing account parameters

Accounts with the admin role can change information for accounts that have lesser permissions. Accounts with the user role cannot.

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
userconfig --change username [-r rolename] [-d description] [-e yes | no]
```

where:

<i>username</i>	Changes the account attribute for <i>username</i> . The account must already exist.
<i>-r rolename</i>	<p>Is an optional argument that changes the role: either <i>admin</i>, <i>switchAdmin</i>, or <i>user</i> in nonsecure mode; <i>admin</i>, <i>user</i>, or <i>nonfcsadmin</i> in secure mode.</p> <p>An account cannot change its own role.</p> <p>You can change the role name of a user-defined account only with a lower level of authorization.</p>
<i>-d description</i>	<p>Is an optional argument; the account description. The description field can be up to 40 printable ASCII characters. The following characters are not allowed: asterisk (*), quotation mark ("), exclamation point (!), semicolon (;), and colon (:).</p> <p>You can change the description of a user-defined account only with a lower level of authorization.</p>
<i>-e</i>	Is an optional argument; enter <i>yes</i> to enable the account or enter <i>no</i> to disable it. If you disable an account, all active CLI sessions for that account are logged out. You can enable or disable user-defined or default accounts.

Recovering user-defined accounts

If a backup account exists (in secure mode), you can recover it with the following command:

```
userConfig --recover
```

The following conditions apply to recovering user accounts:

- Only accounts with admin or higher roles can recover accounts.
- The attributes in the backup database replace the attributes in the current account database.
- An event is stored in the system message log, indicating that accounts have been recovered.

Changing an account password

At each level of account access, you can change passwords for that account and accounts that have lesser privileges.

If you log in to a user account, you can change only that account's password.

If you log in to an admin account, you can change admin and user passwords. You must provide the old password when the account being changed has the same or higher privileges than the current login account. For example, if you are logged in as admin, you need admin passwords to change passwords for admin accounts (except when you change the default user account password at login), but you do not need user passwords to change passwords for user accounts.

A new password must have at least one character different from the old password. The following rules also apply to passwords:

- You cannot change passwords using SNMP.
- Password prompting is disabled when security mode is enabled.
- With Fabric OS 4.4.0 and later, you can use Advanced Web Tools to change admin-level account passwords.
- With Fabric OS 3.2.0 and later, you cannot change default account names.

For information on password behavior when you upgrade (or downgrade) firmware, see ["Effects of firmware changes on accounts and passwords"](#) on page 79.

Changing the password for the current login account

1. Connect to the switch and log in as either admin or user.
2. Issue the password command:
`passwd`
3. Enter the requested information at the prompts.

Changing the password for a different account

1. Connect to the switch and log in as admin.
2. Issue the following password command:
`passwd name`
where *name* is the name of the account.
3. Enter the requested information at the prompts.

If the named account has lesser privileges than the current login account, the old password of the named account is not required. If the named account has equal or higher privileges than the current login account, you are prompted to enter the old password of the named account.

Setting up RADIUS AAA service

Fabric OS 3.2, 4.4.0 and later support RADIUS authentication, authorization, and accounting service (AAA). When configured for RADIUS, a switch becomes a RADIUS client. In this configuration, authentication records are stored in the RADIUS host server database. Login and logout account name, assigned role, and time-accounting records are also stored on the RADIUS server for each user.

By default, RADIUS service is disabled, so AAA services default to the switch local database.

To enable RADIUS service, HP recommends that you access the CLI through an SSH connection so that the shared secret is protected. Multiple login sessions can configure simultaneously; the last session to apply a change leaves its configuration in effect. After a configuration is applied, it persists after a reboot or an HA failover.

The configuration is chassis-based, so it applies to all logical switches (domains) on the switch and replicates itself on a standby CP blade, if one is present. It is saved in a configuration upload and applied in a configuration download.

Configure at least two RADIUS servers so that if one fails, the other assumes service. You can set the configuration with both RADIUS service and local authentication enabled so that if all RADIUS servers do not respond (because of power failure or network problems), the switch uses local authentication.

Consider the following effects of the use of RADIUS service on other Fabric OS features:

- When RADIUS service is enabled, all account passwords must be managed on the RADIUS server. The Fabric OS mechanisms for changing switch passwords remain functional; however, such changes affect only the involved switches locally. They do not propagate to the RADIUS server, nor do they affect any account on the RADIUS server.

When RADIUS is set up for a fabric that contains a mix of switches with and without RADIUS support, the way a switch authenticates users depends on whether a RADIUS server is set up for that switch. For a switch with RADIUS support and configuration, authentication bypasses the local password database. For a switch without RADIUS support or configuration, authentication uses the switch's local account names and passwords.

- When Secure Fabric OS secure mode is enabled, the following behaviors apply:
 - Account passwords stored in the switch-local password database are distributed among all switches in the same fabric. RADIUS configuration is not affected.
 - There are separate admin and nonfcsadmin roles in secure mode. A nonfcsadmin account on a RADIUS server cannot access FCS switches, even if the account is properly authenticated.
 - If a nonfcsadmin account on a RADIUS server logs in to a switch in nonsecure mode, the switch grants the user admin role privileges.
- The following behaviors apply to Advanced Web Tools:
 - Advanced Web Tools client and server keep a session open after a user is authenticated. A password change on a switch invalidates an open session and requires the user to log in again. When integrated with RADIUS, a switch password change on the RADIUS server does not invalidate an existing open session, although a password change on the local switch does.
 - If you cannot log in because of a RADIUS server connection problem, Advanced Web Tools displays a message indicating server outage.

Configuring the RADIUS server

You must know the switch IP address or name to connect to switches. Use the `ipAddrShow` command to display a switch IP address.

For HP StorageWorks SAN Directors (chassis-based systems), the switch IP addresses are aliases of the physical Ethernet interfaces on the CP blades. When specifying client IP addresses for the logical switches in such systems, use the CP blade IP addresses. For accessing both the active and standby CP blade, and for the purpose of HA failover, both of the CP blade IP addresses should be included in the RADIUS server configuration.

User accounts should be set up by their true network-wide identity, rather than by the account names created on a Fabric OS switch. Along with each account name, assign appropriate switch access roles. To manage a nonsecure fabric, these roles can be user or admin. To manage a secure fabric, these roles can be user, admin, or nonfcsadmin.

When they log in to a switch configured with RADIUS, users enter their assigned RADIUS account names and passwords at the prompt. After RADIUS server authenticates a user, it responds with the assigned switch role in an HP Vendor-Specific Attribute (VSA), as defined in the RFC. An Authentication-Accept response without such VSA role assignment, assigns the user role.

The following sections describe how to configure a RADIUS server to support HP clients under different operating systems.

The following procedures work for FreeRADIUS on Solaris and Red Hat Linux. FreeRADIUS is a freeware RADIUS server that you can find at the following web site: www.freeradius.org.

Follow the installation instructions at the web site. FreeRADIUS runs on Linux (all versions), FreeBSD, NetBSD, and Solaris. If you make a change to any of the files used in this configuration, you must stop the server and restart it for the changes to take effect.

FreeRADIUS installation places the configuration files in `$PREFIX/etc/raddb`. By default, the `PREFIX` is `/usr/local`.

Configuring RADIUS service on Linux consists of the following tasks:

- Adding the HP attribute to the server
- Creating the user
- Enabling clients

Adding the attribute to the server

1. Create and save the file `$PREFIX/etc/raddb/dictionary.brocade` with the following information:

```
#
# Brocade FabricOS v5.0.1 dictionary
#
VENDOR      Brocade      1588
#
# attribute 1 defined to be Brocade-Auth-Role
# string defined in user configuration
#
ATTRIBUTE Brocade-Auth-Role 1      string      Brocade
```

This defines the vendor ID as 1588, the vendor attribute 1 as Brocade-Auth-Role, and it is a string value.

2. Open the file `$PREFIX/etc/raddb/dictionary` in a text editor and add the following line:
`$INCLUDE dictionary.brocade`

As a result, the file `dictionary.brocade` is located in the RADIUS configuration directory and loaded for use by the RADIUS server.

Creating the user

Open the `$PREFIX/etc/raddb/user` file in a text editor and add user names and roles for users who will be accessing the switch and authenticating RADIUS. The user logs in using the role specified with Brocade-Auth-Role. The valid roles include root, factory, admin, switchAdmin, and user. You must use quotation marks around "password" and "role".

For example, to set up an account called JohnDoe with the admin role:

```
JohnDoe Auth-Type := Local, User-Password == "johnPassword" Brocade-Auth-Role = "admin"
```

The next example uses the local system password file to authenticate users. (This does not work when using NIS for authentication. The only way to enable authentication with the password file is to force the HP StorageWorks switch to authenticate using PAP; this requires the `-a pap` option with the `aaaConfig` command.) For example:

```
JohnDoe Auth-Type := System, Brocade-Auth-Role = "admin"
```

Enabling clients

Clients are the switches that use the RADIUS server; each client must be defined. By default, all IP addresses are blocked.

On dual-CP switches (Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director), the switch sends its RADIUS request using the IP address of the active CP. When adding clients, add both the active and standby CP IP addresses so that users can still log in, in case of a failover.

1. Open the `$PREFIX/etc/raddb/client.config` file in a text editor and add the switches that are to be configured as RADIUS clients.

For example, to configure the switch at IP address 10.32.170.59 as a client:

```
client 10.32.170.59
    secret      = Secret
    shortname   = Testing Switch
    nastype     = other
```


In this example, `shortname` is an alias used to easily identify the client and `Secret` is the shared secret between the client and server. Make sure that the shared secret matches that configured on the switch (see “[Adding a RADIUS server to the switch configuration](#)” on page 51).

2. Save the file `$PREFIX/etc/raddb/client.config` and then start the RADIUS server as follows:
`$PREFIX/sbin/radiusd`

Windows 2000

Configuring RADIUS service on Windows 2000 consists of the following tasks:

- Installing Internet Authentication Service (IAS). For more information and instructions on installing IAS, see the Microsoft® web site.
- Enabling the Challenge Handshake Authentication Protocol (CHAP). If CHAP authentication is required, Windows must be configured to store passwords with reversible encryption. Reverse password encryption is not the default behavior; it must be enabled.

 **NOTE:** If a user is configured prior to enabling reverse password encryption, the user's password is stored and cannot use CHAP. To use CHAP, the password must be reentered after encryption is enabled. If the password is not reentered, CHAP authentication does not work and the user is then unable to authenticate from the switch.

- Configuring a user: IAS is the Microsoft implementation of a RADIUS server and proxy. IAS uses the Windows native user database to verify user login credentials; it does not list specific users, but instead lists user groups. Each user group should be associated with a specific switch login role. For example, configure a user group for root, admin, factory, switchAdmin, and user, and then add any users whose logins you want to associate to the appropriate group.

Enabling CHAP

1. From the Windows Start menu, select **Programs > Administrative Tools > Local Security Policy** to open the Local Security Settings window.
2. In the Local Security Settings window, expand the **Account Policies** folder and select the **Password Policy** folder.
3. From the list of policies in the Password Policy folder, right-click **Store password using reversible encryption for all users in the domain**, and select **Security** from the pop-up menu.
An additional Local Security Settings window opens.
4. Select the **Enabled** radio button and then click **OK**.


Configuring users

1. From the Windows Start menu, select **Programs > Administrative Tools > Computer Management** to open the Computer Management window.
2. In the Computer Management window, expand the **Local Users and Groups** folder and select the **Groups** folder.
3. Right-click the **Groups** folder and select **New Group** from the pop-up menu.
4. In the New Group window, provide a Name and Description for the group and click **Add**.
5. In the Select Users or Groups window, select the user—who should already have been configured—you want to add to the group and click **Add**.
6. Repeat this for every user you want to add.
7. When you have completed adding all users, click **OK**.
8. In the New Group window, verify that the users you added in [step 4](#) appear in the Members field and then click **Create** to create this group.

The new groups are created for each login type (admin, switchAdmin, user).

Configuring the RADIUS server

1. From the Windows Start menu, select **Programs > Administrative Tools > Internet Authentication Service** to open the Internet Authentication Service window.
2. In the Internet Authentication Service window, right-click the **Clients** folder and select **New Client** from the pop-up menu.


 **NOTE:** A *client* is the device that uses the RADIUS server; in this case, it is the switch.

3. In the Add Client window, provide the following:
 - Friendly name: The friendly name should be an alias that is easily recognizable as the switch to which you are connecting.
 - Protocol: Select **RADIUS** as the protocol.
4. In the Add RADIUS Client window, provide the following:
 - Client address (IP or DNS): Enter the IP address of the switch.
 - Client-Vendor: Select **RADIUS Standard**.
 - Shared secret: Provide a password. Shared secret is a password used between the client device and server to prevent IP address spoofing by unwanted clients. Keep your shared secret password in a safe place. You must enter this password in the switch configuration.
5. Click **Finish** and repeat [step 2](#) through [step 4](#) for all switches on which RADIUS authentication is to be used.
6. In the Internet Authentication Service window, right-click the **Remote Access Policies** folder, and then select **New Remote Access Policy** from the pop-up window.
7. A remote access policy must be created for each login role (root, admin, factory, switchAdmin, and user) for which you want to use RADIUS, so apply this policy to the user groups that you already created.
8. In the Add Remote Access Policy window, enter an easily identifiable **Policy friendly name** that enables you to see the switch login for which the policy is being created, and then click **Next**.
9. After the Add Remote Access Policy window refreshes, click **Add**.
10. In the Select Attribute window, select **Windows Groups** and click **Add**.
11. In the Groups window, click **Add**.
12. In the Select Groups window, select the user-defined group for which you are creating a policy and click **Add**.
13. After adding all appropriate groups, click **OK**.
14. In the Groups window, click **OK**.

15. In the Add Remote Access Policy window, confirm that the **Conditions** section displays the groups that you selected and click **Next**.
16. After the Add Remote Access Policy window refreshes, select the **Grant remote access permission** radio button and click **Next**.
17. After the Add Remote Access Policy window refreshes again, click **Edit Profile**.
18. In the Edit Dial-in Profile window, select the **Authentication** tab and then select only the **Encrypted Authentication (CHAP)** and **Unencrypted Authentication (PAP, SPAP)** check boxes.
19. Select the **Advanced** tab and click **Add**.
20. In the Add Attributes window, select **Vendor-Specific** and click **Add**.
21. In the Multivalued Attribute Information window, click **Add**.
22. In the VSA Information window, select the **Enter Vendor Code** radio button and enter the value 1588.
23. Select the **Yes. It conforms** radio button, and then click **Configure Attribute**.
24. In the Configure VSA (RFC compliant) window, enter the following:
 - a. For the vendor-assigned attribute number, enter the value 1.
 - b. For the attribute format, enter `String`.
 - c. For the attribute value, enter the login role (`root`, `admin`, `factory`, `switchAdmin`, or `user`) the user group must use to log in to the switch.
 - d. Click **OK**.
25. In the Multivalued Attribute Information window, click **OK**.
26. In the Edit Dial-in Profile window, remove all additional parameters (except the one you just added, `Vendor-Specific`) and click **OK**.
27. In the Add Remote Access Policy window, click **Finish**.
28. After returning to the Internet Authentication Service window, repeat [step 6](#) through [step 27](#) to add additional policies for all login types you want to use the RADIUS server. After this is done, you can configure the switch.

Configuring the switch

RADIUS configuration of the switch is controlled by the `aaaConfig` command.

 **NOTE:** On dual-CP switches (Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director), the switch sends its RADIUS request using the IP address of the active CP. When adding clients, add both the active and standby CP IP addresses so that users can still log in to the event of a failover.

The following procedures show how to use the `aaaConfig` command to set up a switch for RADIUS service.

Displaying the current RADIUS configuration

1. Connect to the switch and log in as `admin`.
2. Issue the following command:

```
switch:admin> aaaConfig --show
```

If a configuration exists, its parameters are displayed. If RADIUS service is not configured, only the parameter heading line is displayed. Parameters include:

- **Position:** The order in which servers are contacted to provide service
- **Server:** The server names or IP addresses
- **Port:** The server ports
- **Secret:** The shared secrets
- **Timeouts:** The length of time servers have to respond before the next server is contacted
- **Authentication:** The type of authentication being used on servers

Adding a RADIUS server to the switch configuration

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --add server [-p port] [-s secret] [-t timeout]
[-a pap | chap]
```

where:

<i>server</i>	Is either a server name or an IP address. Avoid duplicating server listings (that is, listing the same server once by name and again by IP address). Up to five servers can be added to the configuration.
<i>-p port</i>	Is an optional argument; enter a server port. The default is port 1812.
<i>-s secret</i>	Is an optional argument; enter a shared secret. The default is <code>sharedsecret</code> . Secrets can be 8 to 40 alphanumeric characters. Make sure that the secret matches that configured on the server.
<i>-t timeout</i>	Is an optional argument; enter the length of time (in seconds) that the server has to respond before the next server is contacted. The default is 3 seconds. Timeout values can range from 1 to 30 seconds.
<i>-a[pap chap]</i>	Specifies PAP or CHAP as the authentication protocol.

Enabling or disabling RADIUS service

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --radius on | off
```

Specifying `on` enables the service; specifying `off` disables it.

At least one RADIUS server must be configured before you can enable RADIUS service.

If no RADIUS configuration exists, turning it on triggers an error message. When the command succeeds, the event log indicates that the configuration is enabled or disabled.

Deleting a RADIUS server from the configuration

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --remove server | all
```

where *server* is a list of servers by either name or IP address. Enter either the name or IP address of the server to be removed.

3. At the prompt, enter `y` to complete the command.

When the command succeeds, the event log indicates that the server is removed.

Changing a RADIUS server configuration

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --change server [-p port] [-s secret] [-t timeout]
[-a pap | chap]
```

where:

<i>server</i>	Is a list of servers by either name or IP address. Enter either the name or IP address of the server to be changed.
---------------	---

- p *port* Is an optional argument; enter a server port.
- s *secret* Is an optional argument; enter a shared secret.
- t *timeout* Is an optional argument; enter the length of time (in seconds) the server has to respond before the next server is contacted.
- a[pap|chap] Specifies PAP or CHAP as authentication protocol.

Changing the order in which RADIUS servers are contacted for service

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> aaaConfig --move server to_position
```

where:

server Is a list of servers by either name or IP address. Enter either the name or IP address of the server whose position is to be changed.

to_position Is the position number to which the server is to be moved.

When the command succeeds, the event log indicates that a server configuration is changed.

Enabling and disabling local authentication

It is useful to enable local authentication so that the switch can take over authentication locally if the RADIUS servers fail to respond because of power outage or network problems. To enable or disable local authentication, issue the following command:

```
switch:admin> aaaConfig --switchdb on | off
```

Specifying *on* enables local authentication; specifying *off* disables it.

When local authentication is enabled and RADIUS servers fail to respond, you can log in to the default switch accounts (admin and user) or any user-defined account. You must know the passwords of these accounts.

RADIUS authentication must be enabled when local database authentication is turned off from the on state; otherwise, an error is returned.

Because local database authentication might be disabled or enabled when enabling or disabling RADIUS authentication, set the local database authentication explicitly to enabled or disabled after setting the desired RADIUS authentication configuration.

When the command succeeds, the event log indicates that local database authentication is disabled or enabled.

Configuring for the SSL protocol

Fabric OS 4.4.0 and later support SSL protocol, which provides secure access to a fabric through Web-based management tools like Advanced Web Tools. SSL support is a standard Fabric OS feature; it is independent of Secure Fabric OS, which requires a license and separate certification.

Switches configured for SSL grant access to management tools through hypertext transfer protocol-secure links (which begin with `https://`) instead of standard links (which begin with `http://`).

SSL uses public key infrastructure (PKI) encryption to protect data transferred over SSL connections. PKI is based on digital certificates obtained from an Internet Certificate Authority (CA), which acts as the trusted key agent.

Certificates are based on the switch IP address or fully-qualified domain name (FQDN), depending on the issuing CA. If you change a switch IP address or FQDN after activating an associated certificate, you might have to obtain and install a new certificate. Check with the CA to verify this possibility, and plan these types of changes accordingly.

Browser and Java support

Fabric OS supports the following web browsers for SSL connections:

- Internet Explorer (Microsoft Windows)
- Mozilla (Solaris and Red Hat Linux)

In countries that allow the use of 128-bit encryption, use the latest version of your browser. For example, Internet Explorer 6.0 and later supports 128-bit encryption by default. You can display the encryption support (called *cipher strength*) using the Internet Explorer Help>About menu option. If you are running an earlier version of Internet Explorer, you might be able to download an encryption patch from the Microsoft web site: <http://www.microsoft.com>.

HP recommends that you upgrade to the Java™ 1.4.2_03 Plug-in on your management workstation. To find the Java version that is currently running, open the Java console and look at the first line of the window.

For details on levels of browser and Java support, see the *HP StorageWorks Fabric OS 5.x Advanced Web Tools administrator guide*.

Summary of SSL procedures

Configure for SSL by obtaining, installing, and activating digital certificates for SSL support. Certificates are required on all switches that are to be accessed through SSL.

You also need to install a certificate to the Java Plug-in on the management workstation, and you might need to add a certificate to your web browser.

Configuring for SSL involves these major steps, which are shown in detail in the next sections:

1. Choose a CA.
2. On each switch:
 - a. Generate a public/private key (`secCertUtil genkey` command).
 - b. Generate a certificate signing request (CSR) (`secCertUtil genscr` command) and store the CSR on an FTP server (`secCertUtil export` command).
3. Obtain the certificates from the CA.

You can request a certificate from a CA through a web browser. After you request a certificate, the CA either sends certificate files by e-mail (public) or provides access to them on a remote host (private). Typically, the CA provides the certificate files listed in [Table 9](#).

Table 9 SSL certificate files

Certificate file	Description
<code>name.crt</code>	The switch certificate.
<code>nameRoot.crt</code>	The root certificate. Typically, this certificate is already installed in the browser, but if not, you must install it.
<code>nameCA.crt</code>	The CA certificate. It is not necessary to install this, but you can if you want the CA name to be displayed in the browser window.

4. On each switch install and activate the certificate.
5. If necessary, install the root certificate to the browser on the management workstation.
6. Add the root certificate to the Java Plug-in keystore on the management workstation.

Choosing a CA

To ease maintenance and allow secure out-of-band communication between switches, consider using one CA to sign all management certificates for a fabric. If you use different CAs, management services operate correctly, but the Advanced Web Tools Fabric Events button is unable to retrieve events for the entire fabric.

Each CA (for example, Verisign or GeoTrust) has slightly different requirements; for example, some generate certificates based on IP address, while others require an FQDN, and most require a 1024-bit public/private key while some might accept a 2048-bit key. Consider your fabric configuration, check CA web sites for requirements, and gather all the information that the CA requires.

Generating a public/private key

Perform the following procedure on each switch:

1. Connect to the switch and log in as admin.
2. Issue the following command to generate a public/private key pair:

```
switch:admin> seccertutil genkey
```

The system reports that this process disables secure protocols, deletes any existing CSR, and deletes any existing certificates.

3. Respond to the prompts to continue and select the key size. For example:

```
Continue (yes, y, no, n): [no] y
Select key size [1024 or 2048]: 1024
Generating new rsa public/private key pair
Done.
```

Because CA support for the 2048-bit key size is limited, select 1024 in most cases.

Generating and storing a CSR

After generating a public/private key (see ["Generating a public/private key"](#) on page 54), perform this procedure on each switch:

1. Connect to the switch and log in as admin.
 2. Issue the following command:
- ```
switch:admin> seccertutil genscr
```
3. Enter the requested information. For example:

```
Country Name (2 letter code, eg, US):US
State or Province Name (full name, eg, California):California
Locality Name (eg, city name):San Jose
Organization Name (eg, company name):Brocade
Organizational Unit Name (eg, department name):Eng
Common Name (Fully qualified Domain Name, or IP address): 192.1.2.3
Generating CSR, file name is: 192.1.2.3.csr
Done.
```

Your CA might require specific codes for Country, State or Province, Locality, Organization, and Organizational Unit names. Make sure that your spelling is correct and matches the CA requirements. If the CA requires that the Common Name be specified as an FQDN, make sure that the FQDN is set on the domain name server.

4. Issue the following command to store the CSR:
- ```
switch:admin> seccertutil export
```

5. Enter the requested information. For example:

```
Select protocol [ftp or scp]: ftp
Enter IP address: 192.1.2.3
Enter remote directory: path_to_remote_directory
Enter Login Name: your account
Enter Password: your password
Success: exported CSR.
```

6. If you are set up for secure file copy protocol, you can select it; otherwise, select `ftp`.
7. Enter the IP address of the switch on which you generated the CSR.
8. Enter the remote directory name of the FTP server to which the CSR is to be sent.
9. Enter your account name and password on the server.

Obtaining certificates

Check the instructions on the CA web site and then perform this procedure for each switch:

1. Generate and store the CSR as described in "[Generating and storing a CSR](#)" on page 54.
2. Open a web browser window on the management workstation and go to the CA web site. Follow the instructions to request a certificate. Locate the area in the request form into which you are to paste the CSR.
3. Through a telnet window, connect to the switch and log in as admin.
4. Issue the following command:

```
switch:admin> seccertutil showcsr
```

The contents of the CSR is displayed.
5. Locate the section that begins with `BEGIN CERTIFICATE REQUEST` and ends with `END CERTIFICATE REQUEST`.
6. Copy and paste this section (including the `BEGIN` and `END` lines) into the area provided in the request form, and then follow the instructions to complete and send the request.

It might take several days to receive the certificates. If the certificates arrive by e-mail, save them to an FTP server. If the CA provides access to the certificates on an FTP server, make note of the path name and make sure you have a login name and password on the server.

Installing a switch certificate

Perform this procedure on each switch:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
switch:admin> seccertutil import
```
3. Select a protocol, enter the IP address of the host on which the switch certificate is saved, and enter your login name and password. For example:

```
Select protocol [ftp or scp]: ftp
Enter IP address: 192.10.11.12
Enter remote directory: path_to_remote_directory
Enter certificate name (must have ".crt" suffix):192.1.2.3.crt
Enter Login Name: your_account
Enter Password: *****
Success: imported certificate [192.1.2.3.crt].
To use this certificate, run the configure command to activate it
```

The certificate is downloaded to the switch.

Activating a switch certificate

Issue the `configure` command and respond to the prompts that apply to SSL certificates:

SSL attributes	Enter yes.
Certificate File	Enter the name of the switch certificate file, for example, 192.1.2.3.crt.
CA Certificate File	If you want the CA name to be displayed in the browser window, enter the name of the CA certificate file; otherwise, skip this prompt.
Select length of crypto key	Enter the encryption key length (40, 56, or 128).
HTTP attributes	Enter yes.
Secure HTTP enabled	Enter yes.

For example:

```
Configure...
System services (yes, y, no, n): [no]
  ssl attributes (yes, y, no, n): [no] yes
Certificate File. (filename or none): [10.33.13.182.crt] 192.1.2.3.crt
  CA Certificate File. (filename or none): [none]
  Select length of crypto key.
    (Valid values are 40, 56, and 128.): (40..128) [128]
http attributes (yes, y, no, n): [no] yes
HTTP Enabled (yes, y, no, n): [yes] no
  Secure HTTP Enabled (yes, y, no, n): [no] yes
```

After you exit the `configure` command, the HTTP daemon restarts to handle HTTPS requests.

Configuring the browser

The root certificate might already be installed on your browser, but if not, you must install it. To determine whether it is already installed, check the certificate store on your browser.

The following procedures are guides for installing root certificates to Internet Explorer and Mozilla browsers. For detailed instructions, see the documentation that came with the certificate.

Checking and installing root certificates on Internet Explorer

1. From the browser **Tools** menu, select **Internet Options**.
2. Click the **Content** tab.
3. Click **Certificates**.
4. Select the various tabs and scroll the lists to see whether the root certificate is listed. If it is listed, you do not need to install it, and you can omit the remainder of this procedure.
5. If the certificate is not listed, click **Import**.
6. Follow the instructions in the Certificate Import wizard to import the certificate.

Checking and installing root certificates on Mozilla

1. From the browser **Edit** menu, select **Preferences**.
2. In the left pane of the Preferences window, expand the **Privacy & Security** list and select **Certificates**.
3. In the right pane, click **Manage Certificates**.
4. In the next window, select the **Authorities** tab.
5. Scroll the authorities list to determine whether the root certificate is listed. (For example, its name might have the form `nameRoot.crt`.) If it is listed, you do not need to install it; omit the remainder of this procedure.
6. If the certificate is not listed, click **Import**.

7. Browse to the certificate location and select the certificate.
For example, select `nameRoot.crt`.
8. Click **Open** and follow the instructions to import the certificate.

Installing a root certificate to the Java Plug-in

For information on Java requirements, see "[Browser and Java support](#)" on page 53.

This procedure is a guide for installing a root certificate to the Java Plug-in on the management workstation. Install the root certificate, if it is not already installed to the plug-in. For detailed instructions, see the documentation that came with the certificate and to the Sun Microsystems web site:

www.sun.com.

1. Copy the root certificate file from its location on the FTP server to the Java Plug-in bin.
For example, the bin location might be:
C: \program files\java\j2re1.4.2_03\bin
2. Open a Command Prompt window and change to the Java Plug-in bin directory.
3. Issue the `keytool` command and respond to the prompts. For example:

```
C:\Program Files\Java\j2re1.4.2_03\bin> keytool -import -alias RootCert -file
RootCert.crt -keystore ..\lib\security\RootCerts
Enter keystore password: changeit
Owner: CN=Brocade, OU=Software, O=Brocade Communications, L=San Jose,
ST=California, C=US
Issuer: CN=Brocade, OU=Software, O=Brocade Communications, L=San Jose,
ST=California, C=US
Serial number: 0
Valid from: Thu Jan 15 16:27:03 PST 2004 until: Sat Feb 14 16:27:03 PST 2004
Certificate fingerprints:
    MD5: 71:E9:27:44:01:30:48:CC:09:4D:11:80:9D:DE:A5:E3
    SHA1: 06:46:C5:A5:C8:6C:93:9C:FE:6A:C0:EC:66:E9:51:C2:DB:E6:4F:A1
Trust this certificate? [no]: yes
Certificate was added to keystore
```

In the example, `changeit` is the default password and `RootCert` is an example of a root certificate name.

Displaying and deleting certificates

[Table 10](#) summarizes the commands that display and delete certificates. For details on these commands, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

Table 10 Commands to display and delete SSL certificates

Command	Description
<code>secCertUtil show</code>	Displays the state of the SSL key and a list of installed certificates
<code>secCertUtil show filename</code>	Displays the contents of a specific certificate
<code>secCertUtil showcsr</code>	Displays the contents of a CSR
<code>secCertUtil delete filename</code>	Deletes a specified certificate
<code>secCertUtil delcsr</code>	Deletes a CSR

Troubleshooting certificates

If you receive messages in the browser or in a pop-up window when logging in to the target switch using HTTPS, see [Table 11](#).

Table 11 SSL messages and actions

Message	Action
The page cannot be displayed	The SSL certificate is not installed correctly or HTTPS is not enabled correctly. Make sure that the certificate has not expired, that HTTPS is enabled, and that certificate file names are configured correctly.
The security certificate was issued by a company you have not chosen to trust.	The certificate is not installed in the browser. Install it as described in "Configuring the browser" on page 56.
The security certificate has expired or is not yet valid	Either the certificate file is corrupted or it needs to be updated. Click View Certificate to verify the certificate content. If it is corrupted or out of date, obtain and install a new certificate.
The name on the security certificate is invalid or does not match the name of the site file	The certificate is not installed correctly in the Java Plug-in. Install it as described in "Installing a root certificate to the Java Plug-in" on page 57.
This page contains both secure and nonsecure items. Do you want to display the nonsecure items?	Click No in this pop-up window. The session opens with a closed lock icon on the lower-right corner of the browser, indicating an encrypted connection.

Configuring SNMP agent and traps

You can perform a configuration for the transmission of SNMP information to management stations. SNMPv3 and SNMPv1 are supported.

The configuration process involves configuring the SNMP agent and configuring SNMP traps. The following commands are used in the process:

- The `configure` command sets the security level. You can specify no security, authentication only, or authentication and privacy.
- The `snmpConfig` command configures the SNMP agent and traps for SNMPv3 or SNMPv1 configurations.
- If necessary for backward compatibility, you can use these legacy commands for the configuration of SNMP v1:
 - The `agtCfgShow`, `agtCfgset`, and `agtCfgDefault` commands configure the SNMPv1 agent.
 - The `snmpMibCapSet` command filters at the trap level and the `snmpMibCapShow` command displays the trap filter values.

The SNMP trap configuration specifies the MIB trap elements to be used to send information to the SNMP management station. There are two main MIB trap choices:

- HP-specific MIB trap is associated with the HP-specific StorageWorks MIB (SW-MIB); it monitors HP StorageWorks switches specifically.
- FibreAlliance MIB trap is associated with the FibreAlliance MIB (FA-MIB); it manages SAN switches and devices from any company that complies with FibreAlliance specifications.

If you use both SW-MIB and FA-MIB, you might receive duplicate information. You can disable the FA-MIB, but not the SW-MIB.

You can also use the following MIBs and their associated traps:

- FICON-MIB (for FICON environments)
- HA-MIB (for the Core Switch 2/64 and SAN Director 2/128)
- SW-EXTTRAP, which includes the Software Serial Number (swSsn) as a part of HP StorageWorks SW traps. It is also used with the legacy Integrated/64 SAN Switch fabrics product to provide detailed group information for a particular trap.

For information on HP StorageWorks MIBs, see the *HP StorageWorks Fabric OS 5.x MIB reference guide*. For information on the specific commands used in these procedures, see online help or the *HP StorageWorks Fabric OS 5.x command reference guide*.

Setting the security level

Use the `configure` command to set the security level (called *SNMP attributes*). You can specify no security, authentication only, or authentication and privacy. For example, to configure for authentication and privacy:

```
switch:admin> configure
```

```
Not all options will be available on an enabled switch.  
To disable the switch, use the "switchDisable" command.
```

```
Configure...
```

```
System services (yes, y, no, n): [no]  
ssl attributes (yes, y, no, n): [no]  
http attributes (yes, y, no, n): [no]  
snmp attributes (yes, y, no, n): [no] y
```

```
Select SNMP Security Level:
```

```
(0 = No security, 1 = Authentication only, 2 = Authentication and  
Privacy): (0..2) [0] 2
```

Using the snmpConfig command

Use the `snmpConfig --set` command to change either the SNMPv3 or SNMPv1 configuration. You can also change access control, MIB capability, and system group.

Sample SNMPv3 configuration:

```
switch:admin> snmpconfig --set snmpv3

SNMPv3 user configuration:
User (rw): [snmpadmin1] adminuser
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 1
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)]: (1..2) [2] 1
New Priv Passwd:
Verify Priv Passwd:
User (rw): [snmpadmin2] shauser
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3] 2
New Auth Passwd:
Verify Auth Passwd:
Priv Protocol [DES(1)/noPriv(2)]: (1..2) [2] 1
New Priv Passwd:
Verify Priv Passwd:
User (rw): [snmpadmin3] nosec
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)]: (2..2) [2]
User (ro): [snmpuser1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv(2)]: (2..2) [2]
User (ro): [snmpuser2]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv(2)]: (2..2) [2]
User (ro): [snmpuser3]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (3..3) [3]
Priv Protocol [DES(1)/noPriv(2)]: (2..2) [2]

SNMPv3 trap recipient configuration:
Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.45.90
UserIndex: (1..6) [1]
Trap recipient Severity level : (0..5) [0] 4
Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.45.92
UserIndex: (1..6) [2]
Trap recipient Severity level : (0..5) [0] 2
Trap Recipient's IP address in dot notation: [0.0.0.0]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Committing configuration...done.
```

Sample SNMPv1 configuration:

```
switch:admin> snmpconfig --set snmpv1

SNMP community and trap recipient configuration:
Community (rw): [Secret C0de] admin
Trap Recipient's IP address in dot notation: [0.0.0.0] 10.32.225.1
Trap recipient Severity level : (0..5) [0] 1
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [10.32.225.2]
Trap recipient Severity level : (0..5) [1]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [10.32.225.3]
Trap recipient Severity level : (0..5) [2]
Community (ro): [public]
Trap Recipient's IP address in dot notation: [10.32.225.4]
Trap recipient Severity level : (0..5) [3]
Community (ro): [common]
Trap Recipient's IP address in dot notation: [10.32.225.5]
Trap recipient Severity level : (0..5) [4]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [10.32.225.6]
Trap recipient Severity level : (0..5) [5]
Committing configuration...done.
```

Sample accessControl configuration:

```
switch:admin> snmpconfig --set accessControl

SNMP access list configuration:
Access host subnet area in dot notation: [0.0.0.0] 192.168.0.0
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0] 10.32.148.0
Read/Write? (true, t, false, f): [true] f
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0] 10.33.0.0
Read/Write? (true, t, false, f): [true] f
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration...done.
```

Sample mibCapability configuration:

```
switch:admin> snmpconfig --show mibCapability
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
    swFabricWatchTrap: YES
    swTrackChangesTrap: NO
FA-TRAP: YES
    connUnitStatusChange: YES
    connUnitEventTrap: NO
    connUnitSensorStatusChange: YES
    connUnitPortStatusChange: YES
SW-EXTTRAP: NO
FICON-TRAP: NO
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: NO
```

Sample systemGroup configuration (default):

```
switch:admin> snmpconfig --default systemGroup
*****
This command will reset the agent's system group configuration back to
factory default
*****
    sysDescr = Fibre Channel Switch
    sysLocation = End User Premise
    sysContact = Field Support
    authTraps = 0 (OFF)

*****
Are you sure? (yes, y, no, n): [no] y
```

Using legacy commands for SNMPv1

Use the `snmpConfig` command to configure the SNMPv1 agent and traps (see ["Using the snmpConfig command"](#) on page 59). However, if necessary for backward compatibility, you can choose to use legacy commands.

Sample SNMP agent configuration information:

```
switch:admin> agtcfshow
Current SNMP Agent Configuration
    Customizable MIB-II system variables:
        sysDescr = FC Switch
        sysLocation = End User Premise
        sysContact = Field Support.
        authTraps = 1 (ON)

SNMPv1 community and trap recipient configuration:
Community 1: Secret C0de (rw)
    Trap recipient: 192.168.1.51
    Trap recipient Severity level: 4
Community 2: OrigEquipMfr (rw)
    Trap recipient: 192.168.1.26
    Trap recipient Severity level: 0
Community 3: private (rw)
    No trap recipient configured yet
Community 4: public (ro)
    No trap recipient configured yet
Community 5: common (ro)
    No trap recipient configured yet
Community 6: FibreChannel (ro)
    No trap recipient configured yet

SNMP access list configuration:
Entry 0: Access host subnet area 192.168.64.0 (rw)]
Entry 1: No access host configured yet
Entry 2: No access host configured yet
Entry 3: No access host configured yet
Entry 4: No access host configured yet
Entry 5: No access host configured yet
```

Sample modification of the SNMP configuration values:

```
switch:admin> agtcfgetset

Customizing MIB-II system variables ...

At each prompt, do one of the followings:
  o <Return> to accept current value,
  o enter the appropriate new value,
  o <Control-D> to skip the rest of configuration, or
  o <Control-C> to cancel any change.

To correct any input mistake:
<Backspace> erases the previous character,
<Control-U> erases the whole line,
sysDescr: [FC Switch]
sysLocation: [End User Premise]
sysContact: [Field Support.]
authTrapsEnabled (true, t, false, f): [true]

SNMP community and trap recipient configuration:
Community (rw): [Secret C0de]
Trap Recipient's IP address in dot notation: [192.168.1.51]
Trap recipient Severity level : (0..5) [0] 3
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address in dot notation: [192.168.1.26]
Trap recipient Severity level : (0..5) [0]
Community (rw): [private]
Trap Recipient's IP address in dot notation: [0.0.0.0] 192.168.64.88
Trap recipient Severity level : (0..5) [0] 1
Community (ro): [public]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [common]
Trap Recipient's IP address in dot notation: [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address in dot notation: [0.0.0.0]

SNMP access list configuration:
Access host subnet area in dot notation: [0.0.0.0] 192.168.64.0
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Access host subnet area in dot notation: [0.0.0.0]
Read/Write? (true, t, false, f): [true]
Committing configuration...done.
value = 1 = 0x1
```

Sample reset of the SNMP agent configuration to default values:

```
switch:admin> agtcfgrdefault
*****
This command will reset the agent's configuration back to factory default
*****

Current SNMP Agent Configuration
Customizable MIB-II system variables:
    sysDescr = Fibre Channel Switch.
    sysLocation = End User Premise
    sysContact = sweng
    authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
    Community 1: Secret C0de (rw)
        Trap recipient: 192.168.15.41
        Trap recipient Severity level: 4
    Community 2: OrigEquipMfr (rw)
        No trap recipient configured yet
    Community 3: private (rw)
        No trap recipient configured yet
    Community 4: public (ro)
        No trap recipient configured yet
    Community 5: common (ro)
        No trap recipient configured yet
    Community 6: FibreChannel (ro)
        No trap recipient configured yet
SNMP access list configuration:
    Entry 0: Access host subnet area 192.168.64.0 (rw)]
    Entry 1: No access host configured yet
    Entry 2: No access host configured yet
    Entry 3: No access host configured yet
    Entry 4: No access host configured yet
    Entry 5: No access host configured yet
*****
Are you sure? (yes, y, no, n): [no] y
Committing configuration...done.
agent configuration reset to factory default
Current SNMP Agent Configuration
Customizable MIB-II system variables:
    sysDescr = Fibre Channel Switch.
    sysLocation = End User Premise
    sysContact = Field Support.
    authTraps = 0 (OFF)
SNMPv1 community and trap recipient configuration:
    Community 1: Secret C0de (rw)
        No trap recipient configured yet
    Community 2: OrigEquipMfr (rw)
        No trap recipient configured yet
    Community 3: private (rw)
        No trap recipient configured yet
    Community 4: public (ro)
        No trap recipient configured yet
    Community 5: common (ro)
        No trap recipient configured yet
    Community 6: FibreChannel (ro)
        No trap recipient configured yet
(output truncated)
```

Sample modification of the options for configuring SNMP MIB traps:

```
switch:admin> snmpmibcapset
The SNMP Mib/Trap Capability has been set to support
FE-MIB
SW-MIB
FA-MIB
FA-TRAP
FA-MIB (yes, y, no, n): [yes]
FICON-MIB (yes, y, no, n): [no] y
HA-MIB (yes, y, no, n): [no] y
SW-TRAP (yes, y, no, n): [no] y
  swFCPortScn (yes, y, no, n): [no]
  swEventTrap (yes, y, no, n): [no]
  swFabricWatchTrap (yes, y, no, n): [no]
  swTrackChangesTrap (yes, y, no, n): [no]
FA-TRAP (yes, y, no, n): [yes]
  connUnitStatusChange (yes, y, no, n): [no]
  connUnitEventTrap (yes, y, no, n): [no]
  connUnitSensorStatusChange (yes, y, no, n): [no]
  connUnitPortStatusChange (yes, y, no, n): [no]
SW-EXTTRAP (yes, y, no, n): [no] y
FICON-TRAP (yes, y, no, n): [no] y
  linkRNIDDeviceRegistration (yes, y, no, n): [no]
  linkRNIDDeviceDeRegistration (yes, y, no, n): [no]
  linkLIRRLListenerAdded (yes, y, no, n): [no]
  linkLIRRLListenerRemoved (yes, y, no, n): [no]
  linkRLIRFailureIncident (yes, y, no, n): [no]
HA-TRAP (yes, y, no, n): [no] y
  fruStatusChanged (yes, y, no, n): [no]
  cpStatusChanged (yes, y, no, n): [no]
  fruHistoryTrap (yes, y, no, n): [no]
Avoid-Duplicate-TRAP (yes, y, no, n): [no] y
switch:admin>
```

The following `snmpMibCapSet` parameters for FA-TRAP appear in the preceding example:

- `connUnitStatusChange`: Indicates that the overall status of the connectivity unit has changed. Its variables are:
 - `connUnitStatus`: The status of the connection unit
 - `connUnitState`: The state of the connection unit
- `connUnitEventTrap`: Indicates that the connectivity unit has generated an event. Its variables are:
 - `connUnitEventId`: The internal event ID
 - `connUnitEventType`: The type of this event
- `connUnitEventObject`: Used with the `connUnitEventType` to identify the object to which the event refers.
- `connUnitEventDescr`: The description of the event.
- `connUnitSensorStatusChange`: Indicates that the status of the sensor associated with the connectivity unit has changed.
- `connUnitSensorStatus`: The status indicated by the sensor.
- `connUnitPortStatusChange`: Indicates that the status of the sensor associated with the connectivity unit has changed.
- `connUnitPortStatus`: Shows overall protocol status for the port.
- `connUnitPortState`: Shows the user-specified state of the port hardware.

Sample view of the SNMP MIB trap setup:

```
switch:admin> snmpmibcapshow
FA-MIB: YES
FICON-MIB: YES
HA-MIB: YES
SW-TRAP: YES
    swFCPortScn: YES
    swEventTrap: YES
    swFabricWatchTrap: YES
    swTrackChangesTrap: YES
FA-TRAP: YES
SW-EXTTRAP: YES
HA-TRAP: YES
    fruStatusChanged: YES
    cpStatusChanged: YES
    fruHistoryTrap: YES
```

Configuring secure file copy

Use the `configure` command to specify that secure file copy (scp) be used for configuration uploads and downloads. For example:

```
switch:admin> configure
```

Not all options will be available on an enabled switch.
To disable the switch, use the "switchDisable" command.

Configure...

```
System services (yes, y, no, n): [no] n
ssl attributes (yes, y, no, n): [no] n
http attributes (yes, y, no, n): [no] n
snmp attributes (yes, y, no, n): [no] n
rpcd attributes (yes, y, no, n): [no] n
cfgload attributes (yes, y, no, n): [no] y
```


```
    Enforce secure config Upload/Download (yes, y, no, n): [no] y
switch:admin>
```

Setting the boot PROM password

The boot PROM password provides an additional layer of security by protecting the boot PROM from unauthorized use. Setting a recovery string for the boot PROM password enables you to recover a lost boot PROM password by contacting your switch service provider. Without the recovery string, a lost boot PROM password cannot be recovered.

Set the boot PROM password and the recovery string on all switches. If your site procedures dictate that you set the boot PROM password without the recovery string, see "[Without a recovery string](#)" on page 69.

To set the boot PROM password with a recovery string, see the section that applies to your switch model.

 **NOTE:** Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. Perform this procedure during planned down time.

4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32

Follow this procedure to set the boot PROM password with a recovery string:

1. Connect to the serial port interface as described in "[Connecting through the serial port](#)" on page 21.
2. Reboot the switch.
3. Press **ESC** within four seconds after the message `Press escape within 4 seconds...` is displayed.

The following options are available:

Option	Description
1 Start system	Continues the system boot process.
2 Recovery password	Lets you set the recovery string and the boot PROM password.
3 Enter command shell	Provides access to boot parameters.

4. Enter 2.

If no password was previously set, the following message is displayed:

```
Recovery password is NOT set. Please set it now.
```

If a password was previously set, the following messages are displayed:

```
Send the following string to Customer Support for password recovery:
```

```
afHTpyLsDolPz0Pk5GzhIw==
```

```
Enter the supplied recovery password.
```

```
Recovery Password:
```

5. Enter the recovery password (string).

The recovery string must be between 8 and 40 alphanumeric characters. HP recommends a random string that is 15 characters or longer for higher security. The firmware prompts for this password only once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

The `New password` prompt is displayed.

6. Enter the boot PROM password, and reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.

The new password is saved (the `saveEnv` command is not required).

7. Reboot the switch.

Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director

The boot PROM and recovery passwords must be set for each CP blade on the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director as follows:

1. Connect to the serial port interface on the standby CP blade, as described in "[Connecting through the serial port](#)" on page 21.
2. Connect to the active CP blade by serial or telnet and issue the `haDisable` command to prevent failover during the remaining steps.
3. For the Core Switch 2/64: Reboot the standby CP blade by pressing the yellow ejector buttons at the top and bottom of the CP blade, and then press both ejector handles back towards the switch to lock the blade back into the slot.

For the SAN Director 2/128 and 4/256 SAN Director: Reboot the standby CP blade by sliding the On/Off switch on the ejector handle of the standby CP blade to Off, and then back to On.

4. Press **ESC** within four seconds after the message `Press escape within 4 seconds...` is displayed.

The following options are available:

Option	Description
1 Start system	Continues the system boot process.
2 Recovery password	Lets you set the recovery string and the boot PROM password.
3 Enter command shell	Provides access to boot parameters.

5. Enter 2.

If no password was previously set, the following message is displayed:

```
Recovery password is NOT set. Please set it now.
```

If a password was previously set, the following messages are displayed:

```
Send the following string to Customer Support for password recovery:
```

```
afHTpyLsDolPz0Pk5GzhIw==
```

```
Enter the supplied recovery password.
```

```
Recovery Password:
```

6. Enter the recovery password (string).

The recovery string must be between 8 and 40 alphanumeric characters. HP recommends a random string that is 15 characters or longer for higher security. The firmware prompts for this password only once. It is not necessary to remember the recovery string because it is displayed the next time you enter the command shell.

The `New password` prompt is displayed.

7. Enter the boot PROM password, and then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.

The new password is saved (the `saveEnv` command is not required).

8. Connect to the active CP blade serially or by telnet and issue the `haEnable` command to restore HA, and then fail over the active CP blade by issuing the `haFailover` command.

Traffic flow through the active CP blade resumes when the failover is complete.


9. Connect the serial cable to the serial port on the new standby CP blade (previously the active CP blade).

10. Repeat [step 2](#) through [step 7](#) for the new standby CP blade (each CP blade has a separate boot PROM password).

11. Connect to the active CP blade serially or by telnet and issue the `haEnable` command to restore high availability.

Without a recovery string

Although you can set the boot PROM password without also setting the recovery string, HP strongly recommends that you set both the password and the string. If your site procedures dictate that you must set the boot PROM password without the string, follow the procedure that applies to your switch model.

 **NOTE:** Setting the boot PROM password requires accessing the boot prompt, which stops traffic flow through the switch until the switch is rebooted. Perform this procedure during planned down time.

4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32


Follow this procedure to set the boot PROM password without a recovery string:

1. Create a serial connection to the switch as described in ["Connecting through the serial port"](#) on page 21.
2. Reboot the switch by issuing the `reboot` command.
3. Press **ESC** within four seconds after the message `Press escape within 4 seconds...` is displayed.

The following options are available:

Option	Description
1 Start system.	Continues the system boot process.
2 Recovery password.	Lets you set the recovery string and the boot PROM password.
3 Enter command shell.	Provides access to boot parameters.

4. Enter 3.
5. At the shell prompt, issue the `passwd` command.

 **NOTE:** The `passwd` command applies only to the boot PROM password when it is entered from the boot interface.

6. Enter the boot PROM password at the prompt, and then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.
7. Issue the `saveEnv` command to save the new password.
8. Reboot the switch by issuing the `reset` command.

Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director

On the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, set the password on the standby CP blade, fail over, and then set the password on the previously active (now standby) CP blade to minimize disruption to the fabric:

1. Determine the active CP blade by opening a telnet session to either CP blade, connecting as admin, and entering the `haShow` command.
2. Connect to the active CP blade serially and by telnet and issue the `haDisable` command to prevent failover during the remaining steps.
3. Create a serial connection to the standby CP blade as described in ["Connecting through the serial port"](#) on page 21.
4. For the Core Switch 2/64: Reboot the standby CP blade by pressing the yellow ejector buttons at top and bottom of the CP blade, and then press both ejector handles back towards the switch to lock the blade back into the slot.

SAN Director 2/128 and 4/256 SAN Director: Reboot the standby CP blade by sliding the On/Off switch on the ejector handle of the standby CP blade to Off, and then back to On.


This causes the blade to reset.

5. Press **ESC** within four seconds after the message `Press escape within 4 seconds...` is displayed.

The following options are available:

Option	Description
1 Start system.	Continues the system boot process.
2 Recovery password.	Lets you set the recovery string and the boot PROM password.
3 Enter command shell.	Provides access to boot parameters.

6. Enter 3.
7. Issue the `passwd` command at the shell prompt.

 **NOTE:** The `passwd` command applies only to the boot PROM password when it is entered from the boot interface.

8. Enter the boot PROM password at the prompt, and then reenter it when prompted. The password must be eight alphanumeric characters (any additional characters are not recorded). Record this password for future use.
9. Issue the `saveEnv` command to save the new password.
10. Reboot the standby CP blade by issuing the `reset` command.
11. Connect to the active CP blade serially and by telnet and issue the `haEnable` command to restore HA, and then fail over the active CP blade by issuing the `haFailover` command.
Traffic resumes flowing through the newly active CP blade after it has completed rebooting.
12. Connect the serial cable to the serial port on the new standby CP blade (previously the active CP blade).
13. Repeat [step 3](#) through [step 10](#) for the new standby CP blade.
14. Connect to the active CP blade serially and by telnet and issue the `haEnable` command to restore HA.

Recovering forgotten passwords

If you know the root password, you can use this procedure to recover the user, admin, and factory passwords:

1. Open a CLI connection (serial or telnet) to the switch. If secure mode is enabled, connect to the primary FCS switch.
2. Log in as root.
3. Issue the command for the type of password that was lost:
 - `passwd user`
 - `passwd admin`
 - `passwd factory`
4. Enter the requested information at the prompts.

To recover a lost root password, contact your switch service provider.

To recover a lost boot PROM password, contact your switch service provider. You must have previously set a recovery string to recover the boot PROM password.

4 Maintaining configurations and firmware

This chapter contains procedures for maintaining switch configurations and maintaining firmware.

Maintaining configurations

It is important to maintain consistent configuration settings on all switches in the same fabric, because inconsistent parameters (such as inconsistent PID formats) can cause fabric segmentation. As part of standard configuration maintenance procedures, HP recommends that you back up all important configuration data for every switch on a host computer server for emergency reference.

The following sections contain procedures for basic switch configuration maintenance.

Displaying configuration settings

The switch configuration file contains four sections:

- The Boot Parameters section contains variables, such as the switch's name and IP address.
- The Licenses section lists the licenses that are active on the switch.
- The Chassis Configuration section contains configuration variables, such as diagnostic settings, fabric configuration settings, and SNMP settings.
- The Configuration section contains licensed option configuration parameters.

To display configuration settings, connect to the switch, log in as admin, and issue the `configShow` command. The configuration settings vary depending on switch model and configuration.

Backing up a configuration

If the configuration is lost or unintentional changes are made, keep a backup copy of the configuration file (or a backup copy of both configuration files, one for each logical switch—if you are using a Core Switch 2/64 or SAN Director 2/128 with two domains). The configuration file is what gets saved when you issue the `configUpload` command. Keep individual backup files for all switches in the fabric. Do not copy configurations from one switch to another.

The following information is not saved in a backup file:

- `dnsconfig` information
- Passwords

You must have a valid account on the FTP server where the backup file is stored.

You can specify the use of secure file copy (`scp`) during the procedure. For instructions on configuring the use of `scp` by default, see "[Configuring secure file copy](#)" on page 67.

Before beginning, verify that you can reach the FTP server from the switch. Using a telnet connection, save a backup copy of the configuration file from a logical switch to a host computer as follows:

1. Verify that the FTP service is running on the host computer.
2. Connect to the switch and log in as admin.
3. Issue the `configUpload` command.

The command becomes interactive and you are prompted for the required information.

4. Respond to the prompts as follows:

Protocol	If your site requires the use of Secure Copy, specify <code>scp</code> . Otherwise, specify <code>ftp</code> .
Server Name or IP Address	Enter the name or IP address of the server where the file is to be stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled. For details about the <code>dnsConfig</code> command, see the <i>HP StorageWorks Fabric OS 5.x command reference guide</i> .
User name	Enter the user name of your account on the server, for example, JohnDoe.
File name	Specify a file name for the backup file, for example, <code>config.txt</code> . Absolute path names can be specified using forward slash (/). Relative path names create the file in the user's home directory on UNIX servers, and in the directory where the FTP server is running on Windows servers.
Password	Enter your account password for the server.

For example:

```
switch:admin> configupload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxxx
Upload complete
switch:admin>
```

Restoring a configuration

Restoring a configuration involves overwriting the configuration on the switch by downloading a previously saved backup configuration file. Perform this procedure during a planned down time.

Make sure that the configuration file you are downloading is compatible with your switch model; configuration files from other model switches might cause your switch to fail.

You must have a user ID on the FTP server where the backup file is stored.

Use the following procedure:

1. Verify that the FTP service is running on the server where the backup configuration file is located.
2. Connect to the switch and log in as admin.
3. Disable the switch by issuing the `switchDisable` command.
4. Issue the `configDownload` command.

The command becomes interactive and you are prompted for the required information.

5. Respond to the prompts as follows:

Protocol	If your site requires the use of Secure Copy, specify <code>scp</code> . Otherwise, specify <code>ftp</code> .
Server Name or IP Address	Enter the name or IP address of the server where the file is stored; for example, 192.1.2.3. You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server, for example, JohnDoe.
File name	Specify the full path name of the backup file, for example, <code>/pub/configurations/config.txt</code> .
Password	Enter your account password for the server.

6. At the Do you want to continue [y/n] prompt, enter y.
7. Wait for the configuration to be restored.
8. When the process is finished, issue the `switchEnable` command.


For example:

```
switch:admin> configdownload
Protocol (scp or ftp) [ftp]: ftp
Server Name or IP Address [host]: 192.1.2.3
User Name [user]: JohnDoe
File Name [config.txt]: /pub/configurations/config.txt
Password: xxxxxx

*** CAUTION ***

This command is used to download a backed-up configuration
for a specific switch. If using a file from a different
switch, this file's configuration settings will override
any current switch settings. Downloading a configuration
file, which was uploaded from a different type of switch,
may cause this switch to fail.

Do you want to continue [y/n]: y
download complete..
switch:admin> switchenable
```

 **NOTE:** After you download a configuration file, you must reboot to be sure the parameters are enabled. Before the reboot, this type of parameter is listed in the configuration file, but it is not effective until after the reboot.

Restoring configurations in a FICON environment

If the switch is operating in a FICON Control Unit Port (CUP) environment, and the ASM (active=saved) bit is set on, the switch ignores the initial program load (IPL) file that is downloaded when you restore a configuration. [Table 12](#) describes this behavior in detail.

Table 12 Backup and restore in a FICON CUP environment

ASM bit	Command	Description
on or off	<code>configupload</code>	All the files saved in the file access facility are uploaded to the management workstation. A section in the uploaded configuration file labeled FICON_CUP is in an encoded format.
on	<code>configdownload</code>	Files saved on the switch that are also present in the FICON_CUP section of the configuration file are overwritten. Files in the FICON section of configuration file that are not currently present on the switch are saved. The IPL file is not replaced, because active=saved mode is on. A message is displayed in the syslog to warn that the IPL file is not being overwritten.
off	<code>configdownload</code>	Files saved on the switch that are also present in the FICON_CUP section of the configuration file are overwritten. Files in the FICON section of configuration file that are not currently present on the switch are saved. The IPL file is replaced, because active=saved mode is off.

If `fmsmode` is enabled in a configuration file, but is disabled on the switch, the `configdownload` command fails and displays an error message. This prevents undesirable conditions that could result from enabling `fmsmode` on a switch that does not require it.

Downloading configurations across a fabric

To save time when configuring fabric parameters and software features, you can save a configuration file from one switch and download it to other switches of the same model type, as shown in the following procedure. Avoid downloading configuration files to different model switches, because that can cause the switches to fail.

1. Configure one switch first.
2. Use the `configUpload` command to save the configuration information.
See “[Backing up a configuration](#)” on page 73.
3. Use the `configDownload` command to download the file onto each of the remaining switches.
See “[Restoring a configuration](#)” on page 74.

Printing hard copies of switch information

HP recommends that you print a hard copy of all key configuration data, including license key information for every switch, and store it in a secure place for emergency reference. Print out the information from the following commands, and store the printouts in a secure location:

- The `configShow` command displays configuration parameters and setup information, including license information.
- The `ipAddrShow` command displays the IP address.
- The `licenseShow` command displays the license keys you have installed and provides better detail than the license information from the `configShow` command.

Depending on the security procedures of your company, you might also want to keep a record of the user levels and passwords (including any boot ROM passwords) for all switches in the fabric. Access to this sensitive information should be limited.

Maintaining firmware

This section explains how to obtain and install firmware. Fabric OS 5.0.1x provides nondisruptive firmware installation.

In most cases, you will be upgrading firmware; that is, installing a later firmware version than the one you are currently running. However, some circumstances might require installing an earlier version; that is, downgrading the firmware. The procedures in this section assume that you are upgrading firmware, but they work for downgrading as well, provided the old and new firmware versions are compatible. Always reference the latest release notes for updates that may exist regarding downgrades under particular circumstances.

Using the CLI (or HP Advanced Web Tools), you can upgrade the firmware on one switch at a time. You can also use the optionally licensed HP Fabric Manager software tool to upgrade firmware simultaneously on multiple switches. For details on Fabric Manager and other licensed software tools, visit the HP web site: <http://h18006.www1.hp.com/storage/saninfrastructure/switches.html>.

Obtaining and unzipping firmware

Firmware upgrades are available for customers with support service contracts and partners on the HP Storage web site: <http://welcome.hp.com/country/us/eng/prodserv/storage.html>. For currently sold switches:

1. Locate the **Networked storage** section under **IT storage products** and click **Storage area networks**.
The SAN Infrastructure page is displayed.
2. Click **Fibre Channel Switches**.
The Fibre Channel switches page is displayed.
3. Go to the **B-Series Fabric-Enterprise Class** section and select the appropriate switch.
The switch overview page is displayed.

4. In the **Product information** section on the right side, select **Software & drivers**.

The download drivers & software page is displayed.

5. Click the appropriate switch in the **select your product** section.

The specify operating system page is displayed.

6. Click **Cross operating system (BIOS, Firmware, Diagnostics, etc.)**.

The download drivers and software page is displayed.

7. In the **Firmware** section, click the blue **download** button to the right of the applicable firmware.

To locate all available switch firmware, start at the HP web site: <http://www.hp.com> and select **Driver Downloads**.

The Software & Driver Downloads page is displayed. You may search for your product using either of the following methods:

1. Select the **Download drivers and software** radio button, enter your product name in the space provided, and press **Enter**.

The Product search results page is displayed.

- a. Select the appropriate product.

The specify operating system page is displayed.

- b. Click **Cross operating system (BIOS, Firmware, Diagnostics, etc.)**.

The download drivers and software page is displayed.

- c. In the **Firmware** section, click the blue **download** button to the right of the applicable firmware.

2. Click **Storage** in the **Or Select a product category** section.

The Storage page is displayed.

- a. Click **SAN Infrastructure**.

The SAN Infrastructure page is displayed.

- b. Select the appropriate product family.

The product family page is displayed.

- c. Select the appropriate switch.

The specify operating system page is displayed.

- d. Click **Cross operating system (BIOS, Firmware, Diagnostics, etc.)**.

The download drivers and software page is displayed.

- e. In the **Firmware** section, click the blue **download** button to the right of the applicable firmware.

Before you can use the `firmwareDownload` command to update the firmware on your equipment, you must unzip the firmware (using the UNIX `tar` or `gzip` command or a Windows unzip program).

When you unpack the downloaded firmware it expands into a directory that is named according to the version of Fabric OS it contains. For example, if you download and unpack `Fabric OS 5.0.1.zip`, it expands into a directory called `5.0.1`. When you use the `firmwaredownload` command, you specify the path to the version `5.0.1` directory and append the keyword `release.plist` to the path.

Checking connected switches

If the switch to be upgraded is running version 4.1.0 firmware or later, HP recommends that all switches directly connected to it be running versions no earlier than 2.6.1, 3.1.0, or 4.1.0. If some connected switches are running earlier firmware versions, upgrade them to at least the earliest recommended version (shown in [Table 13](#)) before upgrading firmware on your switch. HP recommends that you download the latest firmware; to download firmware, see "[Obtaining and unzipping firmware](#)" on page 76.

Table 13 Recommended firmware

Switch model ¹	Earliest recommended Fabric OS version
4/8 SAN Switch and 4/16 SAN Switch	5.0.1
1 GB Switches	2.6.1
SAN Switch 2/8-EL, SAN Switch 2/16-EL, and SAN Switch 2/16	3.1.0
SAN Switch 2/8V and SAN Switch 2/16V	4.2.0
SAN Switch 2/32	4.1.0
Brocade 4Gb SAN Switch for HP p-Class BladeSystem	5.0.1
SAN Switch 4/32	4.4.0
Core Switch 2/64	4.1.0
SAN Director 2/128	4.2.0
4/256 SAN Director	5.0.1
<p>1. During code activation on the SAN Switch 2/8V, SAN Switch 2/16V, or SAN Switch 2/32 running Fabric OS 4.1.0 or later, data continues to flow between hosts and storage devices; however, fabric services are unavailable for a period of approximately 50–55 seconds. Possible disruption of the fabric can be minimized by ensuring that switches logically adjacent to these models (directly connected via an ISL) are running at the minimum Fabric OS 2.6.1 or later, 3.1.0 or later, or 4.1.0 or later.</p> <p>If the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, or SAN Switch 4/32 are adjacent and you start firmware downloads on them at same time, there might be I/O disruption.</p>	

To determine whether you need to upgrade connected switches before upgrading your switch, use the following procedure on each connected switch to display firmware information and build dates.

1. Connect to the switch and log in as admin.
2. Issue the `version` command.

The following information is displayed:

- `Kernel` displays the version of switch kernel operating system.
- `Fabric OS` displays the version of switch Fabric OS.
- `Made on` displays the build date of firmware running in switch.
- `Flash` displays the installation date of firmware stored in nonvolatile memory.
- `BootProm` displays the version of the firmware stored in the boot PROM.

About the download process

The `firmwareDownload` command downloads unzipped switch firmware from an FTP server to the switch's nonvolatile storage area.

In the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, this command (when not using any options flags) by default downloads the firmware image to a standby CP, if there is one, to prevent disruption to application services. This operation depends on HA support. If HA is not available, experienced technicians can upgrade the CPs one at a time, using the `-s` option.

- △ **CAUTION:** To ensure a nondisruptive download, for each nondirector class switch in your fabric, complete all firmware download changes before issuing the `firmwareDownload` command on the next switch.

HP StorageWorks fixed-port models and each CP blade of the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director have two partitions of nonvolatile storage areas (a primary and a secondary) to store two firmware images. The `firmwareDownload` command always loads the new image into the secondary partition and swaps the secondary partition to be the primary. It then reboots the partition and activates the new image. Finally, it performs the `firmwareCommit` procedure, to copy the new image to the other partition.

Effects of firmware changes on accounts and passwords

Table 14 describes what happens to accounts and passwords when you replace the switch firmware with a different version. *Upgrading* means installing a later version of firmware. *Downgrading* means installing an earlier version.

Table 14 Effects of firmware changes on accounts and passwords

Change	First time	Subsequent times (after upgrade, then downgrade, and then upgrade)
Upgrading	Default accounts and their passwords are preserved.	User-defined and default accounts and their passwords are preserved.
Downgrading	User-defined accounts are no longer valid. Default accounts and their passwords are preserved. If a default account was disabled, it is reenabled after the downgrade.	User-defined and default accounts and their passwords are preserved, including accounts added after the first upgrade.
Upgrading to version 3.2.0	(You might upgrade a switch in the fabric as part of “ Checking connected switches ” on page 77.) Earlier versions allowed you to change the default account names. You cannot add user-defined accounts until you change the names back to default with the <code>passwdDefault</code> command.	

Considerations for downgrading firmware

The following items must be considered before attempting to downgrade to an earlier version of Fabric OS:

- If your fabric is set to the extended edge PID format and you want to downgrade to an earlier Fabric OS version that does not support extended edge, you must change the PID to a supported format. For more information, see “[Configuring the PID format](#)” on page 213.
- Downgrading a SAN Director 2/128 that is configured for two domains from Fabric OS 4.4.0 to Fabric OS 4.2.0 is not supported.
- If you are running Fabric OS 4.0.2 firmware on a SAN Switch 2/32, you cannot downgrade to an earlier version.

Considerations for FICON CUP environments

To prevent channel errors during nondisruptive firmware installation, the switch CUP port must be taken offline from all host systems.

Upgrading HP StorageWorks switches

The 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32 maintain primary and secondary partitions for firmware. The `firmwareDownload` command defaults to an autocommit option that copies the firmware from one partition to the other.

Do not override an autocommit under normal circumstances; use the default. See “[Upgrading firmware in single-CP mode](#)” on page 239 for details about overriding the autocommit option.

As an alternative, before starting a firmware download, you can connect the switch with a serial console cable to a computer that is running a session capture. The information collected might be useful for troubleshooting.

Summary of the upgrade process


The following summary describes the default behavior of the `firmwareDownload` command (without options) on the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32:

1. Issue the `firmwareDownload` command.
2. Fabric OS downloads firmware to the secondary partition.
3. The system performs an HA reboot (`haReboot`). After the `haReboot`, the former secondary partition is the primary partition.
4. The system replicates the firmware from the primary to the secondary partition.

You can issue the `firmwareDownloadStatus` command to view the firmware process.

Upgrading 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32

The upgrade process first downloads and then commits the firmware to the switch. While the upgrade is proceeding, you can start another telnet session on the switch and observe the upgrade progress.

 **NOTE:** After you start the process, do not enter any disruptive commands (such as reboot) that interrupt the process. The entire firmware download and commit process takes approximately 17 minutes. If there is a problem, wait for the timeout (30 minutes for network problems; 10 minutes for incorrect IP address). Disrupting the process can render the switch inoperable and require you to seek help from Customer Support.

Do not disconnect the switch from power during the process; the switch could become inoperable upon reboot.

Use this procedure to upgrade firmware for the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32:

1. Verify that the FTP service is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the HP web site:
<http://welcome.hp.com/country/us/en/prodserv/storage.html> and store the file on the FTP server.
Verify that the FTP service is running and unpack the compressed files preserving directory structures.
3. Use the `firmwareShow` command to check the current firmware version on connected switches.
Upgrade their firmware, if necessary, before upgrading this switch.
See “[Checking connected switches](#)” on page 77.
4. Connect to the switch and log in as admin.
5. Use the `firmwareShow` command to check the current firmware version of the switch to verify compatibility with the version of firmware you are going to download.

 **NOTE:** For the SAN Switch 2/8V and SAN Switch 2/16V: If you are running Fabric OS 4.2.0 firmware, you cannot downgrade to earlier versions.

For the SAN Switch 2/32: If you are running Fabric OS 4.0.2 firmware, you cannot downgrade to earlier versions.

6. Issue the `firmwareDownload` command.
7. At the Do you want to continue [y/n] prompt, enter `y`.
8. Respond to the prompts as follows:

Server Name or IP Address	Enter the name or IP address of the server where the firmware file is stored, for example, 192.1.2.3. You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server, for example, JohnDoe.
File name	Specify the full path name of the firmware directory, appended by <code>release.plist</code> , for example, <code>/pub/v5.0.1/release.plist</code> . For version 4.x and 5.x switches only, do not attempt to locate the <code>release.plist</code> file in the top level directory; there is a <code>release.plist</code> file for each platform, and the correct one is selected.
Password	Enter your account password for the server.

After the firmware is downloaded, the switch reboots and starts the firmware commit.

9. After the reboot, connect to the switch and log in again as admin.
10. If you want to watch the upgrade progress, issue the `firmwareDownloadStatus` command.
11. After the firmware commit finishes, issue the `firmwareShow` command to display the firmware level for both partitions.

For example:

```
switch:admin> firmwaredownload
You can run firmwareDownloadStatus to get the status of this command.
This command will cause the switch to reset and will require that existing
telnet, secure telnet or SSH sessions be restarted.
Do you want to continue [Y]: y
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v5.0.1/release.plist
Password: xxxxxx
Firmwaredownload has started.

0x8fd (Fabric OS): Switch: 0, Warning SULIB-FWDL_START, 3, Firmwaredownload
command has started.
.
.
.
```

Log in again to view the upgrade progress:

```
switch:admin> firmwaredownloadstatus
[0]: Tue Apr 20 10:32:34 2004
cp0: Firmwaredownload has started.
[1]: Tue Apr 20 10:36:07 2004
cp0: Firmwaredownload has completed successfully.
[2]: Tue Apr 20 10:57:09 2004
cp0: Firmwarecommit has started.
[3]: Tue Apr 20 10:36:07 2004
cp0: Firmwarecommit has completed successfully.
[4]: Tue Apr 20 11:03:28 2004
cp0: Firmwaredownload command has completed successfully.
switch:admin> firmwareshow
Primary partition: v5.0.1
Secondary Partition: v5.0.1
switch:admin>
```

△ **CAUTION:** To successfully download firmware to a director you must have an active Ethernet connection on *both* CPs.

Upgrading HP StorageWorks directors

You can download firmware to the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director without disrupting the overall fabric if the two CP blades are installed and fully synchronized. Use the `haShow` command to confirm synchronization. If only one CP blade is powered on, the switch must reboot to activate firmware, which is disruptive to the overall fabric.


If there is an error during the firmware download, the system ensures that the two partitions of a CP blade contain the same version of firmware. However, the two CP blades might contain different versions of firmware; in that event, repeat the firmware download process.

During the upgrade process, the director fails over to its standby CP blade and the IP addresses for the two logical switches move to that CP blade's Ethernet port. This might cause informational ARP address reassignment messages to appear on other switches in the fabric. This is normal behavior, because the association between the IP addresses and MAC addresses has changed.

Summary of the upgrade process

The following summary describes the default behavior of the `firmwareDownload` command (without options) on Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director:


1. Issue the `firmwareDownload` command on the active CP blade.
2. The standby CP blade downloads firmware.
3. The standby CP blade reboots and comes up with the new Fabric OS.
4. The active CP blade synchronizes its state with the standby CP blade.
5. The active CP blade forces a failover and reboots to become the standby CP blade.
6. The new standby CP blade (the active CP blade before the failover) downloads firmware.
7. The new standby CP blade reboots and comes up with the new Fabric OS.
8. The new active CP blade synchronizes its state with the new standby CP blade.
9. The `firmwareCommit` command runs on both CP blades.

 **NOTE:** After you start the process, do not issue any disruptive commands (such as reboot) that will interrupt the process. The entire firmware download and commit process takes approximately 15 minutes. If there is a problem, wait for the timeout (30 minutes for network problems; 10 minutes for incorrect IP address). Disrupting the process can render the switch inoperable and require you to seek help from Customer Support.

Do not disconnect the switch from power during the process, because the switch could become inoperable upon reboot.

Upgrading the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director

Core Switch 2/64 directors have four IP addresses: one for each of the two logical switches (switch 0 and switch 1) and one for each of the two CP blades (CP0 in slot 5 and CP1 in slot 6). The SAN Director 2/128 in its default configuration has three IP addresses, but can be configured for four. The 4/256 SAN Director does not support two domains; hence, you can use only three IP addresses.

 **NOTE:** By default, the `firmwareDownload` command upgrades both the active CP blade and the standby CP blade. When upgrading a Core Switch 2/64 that is running 4.0.0c or earlier, you must upgrade each CP blade separately, as described in “[Upgrading a single Core Switch 2/64 or SAN Director 2/128 blade](#)” on page 240. (Do not use the following procedure under normal circumstances.)

Follow this procedure to upgrade the firmware on Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director:

1. Verify that the FTP service is running on the host server and that you have a user ID on that server.
2. Obtain the firmware file from the HP web site:
<http://welcome.hp.com/country/us/en/prodserv/storage.html> and store the file on the FTP server. Verify that the FTP service is running and unpack the compressed files preserving directory structures.
3. Use the `firmwareShow` command to check the current firmware version on connected switches. Upgrade the firmware, if necessary, before proceeding with upgrading this switch.
See “[Checking connected switches](#)” on page 77.
4. Using a telnet session, connect to the switch and log in as admin.
5. For the Core Switch 2/64, use the `firmwareShow` command to check the current firmware version of the switch.
6. Issue the `haShow` command to confirm that the two CP blades are synchronized.

CP blades must be synchronized and running Fabric OS 4.2.0 or later to provide a nondisruptive download. If the two CP blades are not synchronized, and the current firmware version is 4.2.0 or later, issue the `haSyncStart` command to synchronize the two CP blades. In the following example, the active CP blade is CP1 and the standby is CP0:

```
switch:admin> hashow
Local CP (Slot 6, CP1): Active
Remote CP (Slot 5, CP0): Standby
HA Enabled, Heartbeat up, HA State is in Sync
switch:admin>
```

7. Log in to either of the logical switches (sw0 for the 4/256 SAN Director and SAN Director 2/128 configured with a single domain).
8. Issue the `firmwareDownload` command.
9. At the Do you want to continue [y/n] prompt, enter y.

10. Respond to the prompts as follows:

Server Name or IP Address	Enter the name or IP address of the server where the firmware file is stored, for example, 192.1.2.3. You can enter a server name if DNS is enabled.
User name	Enter the user name of your account on the server, for example, JohnDoe.
File name	Specify the full path name of the firmware directory, appended by release.plist, for example, /pub/v5.0.1/release.plist.
Password	Enter your account password for the server.

The firmware is downloaded to one CP blade at a time, beginning with the standby CP blade. During the process, the active CP blade is failed over. After the firmware is downloaded, a firmware commit starts on both CP blades. The entire firmware download and commit process takes approximately 35 minutes.

11. Optional: After the failover, connect to the switch and log in again as admin.

12. Using a separate telnet session, issue the `firmwareDownloadStatus` command to monitor the firmware download status.

13. Issue the `firmwareShow` command to display the new firmware versions.

For example:

```
switch:admin> firmwaredownload
This command will upgrade both CPs in the switch. If you
what to upgrade a single CP only, please use -s option.

You can run firmwareDownloadStatus to get the status
of this command.

This command will cause the active CP to reset and will
require that existing telnet, secure telnet, or SSH sessions
be restarted.


Do you want to continue [Y]: y
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v5.0.1/release.plist
Password:*****
FirmwareDownload has started on Standby CP. It may take up to 30 minutes.
Firmwaredownload has completed successfully on Standby CP.
.
.
.
Standby CP reboots.
Standby CP booted up.
Standby CP booted up with new firmware.
cp1: Firmwarecommit has started on both Active and Standby CPs.
cp1: Firmwarecommit has completed successfully on Active CP.
cp1: Firmwaredownload command has completed successfully.
switch:admin>
```

Start a new session to view the upgrade progress:

```
switch:admin> firmwaredownloadstatus
[0]: Tue Apr 20 15:18:56 2003
cp0: Firmwaredownload has started on Standby CP. It may take up to 10 minutes.
[1]: Tue Apr 20 15:24:17 2003
cp0: Firmwaredownload has completed successfully on Standby CP.
[2]: Tue Apr 20 15:24:19 2003
cp0: Standby CP reboots.
[3]: Tue Apr 20 15:27:06 2003
cp0: Standby CP booted up.
[4]: Tue Apr 20 15:29:01 2003
cp1: Active CP forced failover succeeded. Now this CP becomes Active.
[5]: Tue Apr 20 15:29:05 2003
cp1: Firmwaredownload has started on Standby CP. It may take up to 30 minutes.
[6]: Tue Apr 20 15:34:16 2003
cp1: Firmwaredownload has completed successfully on Standby CP.
[7]: Tue Apr 20 15:34:19 2003
cp1: Standby CP reboots.
[8]: Tue Apr 20 15:36:59 2003
cp1: Standby CP booted up with new firmware.
[9]: Tue Apr 20 15:37:04 2003
cp1: Firmwarecommit has started on both Active and Standby CPs.
[10]: Tue Apr 20 15:42:48 2003
cp1: Firmwarecommit has completed successfully on Active CP.
[11]: Tue Apr 20 15:42:49 2003
cp1: Firmwaredownload command has completed successfully.
```

Troubleshooting firmware downloads

A firmware download can fail for many reasons, such as a power failure, a failed network connection, a failed FTP server, or an incorrect path to unpacked firmware files. In most cases, the firmware is not affected. You can make necessary corrections (for example, check the Ethernet cables and check the file path names) and then rerun the `firmwareDownload` command.

 **NOTE:** Under firmware versions earlier than 4.1.0, do not perform a firmware download while the switch is running POST. If a firmware download is attempted on a Core Switch 2/64 while POST is running, the download might fail because the CP blades cannot synchronize with each other.

Issue the `firmwareShow` command to see whether both CP blades have the same firmware. In the following example, the active and standby CP blades have the same version:

```
switch: admin> firmwareshow
Local CP (Slot 6, CP1): Standby
    Primary partition:      v5.0.1d
    Secondary Partition:    v5.0.1d
Remote CP (Slot 5, CP0): Active
    Primary partition:      v5.0.1d
    Secondary Partition:    v5.0.1d

Note: If Local CP and Remote CP have different versions
of firmware, please retry firmwaredownload command.
switch: admin>
```

Decide which firmware version you want to be applied to each CP blade. If you want the version on the standby CP, issue the `haFailover` command on the active CP. If you want the version from the active CP, issue the `firmwareDownload -s` command on the standby CP. After entering the `haFailover` command, you must issue the `firmwareDownload -s` command on the new standby CP.

5 Configuring Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director

This chapter contains procedures that are specific to the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director.

Because directors contain interchangeable 16-port blades (32-port blades in the 4/256 SAN Director), their procedures differ from those for the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32 fixed-port switches. For example, fixed-port models identify ports by domain,port number, while director models identify ports by slot/port number.

Also, because the Core Switch 2/64 director comprises two logical switches (domains), and the SAN Director 2/128 and 4/256 SAN Director in their default configurations have only one domain (the 4/256 SAN Director supports only one domain), procedures for the directors sometimes differ from one another.

For detailed information about the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, see the installation guide for the switch.

Identifying ports

The Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director have slots and can have a variable number of ports within a given domain. Ports are identified by their combined slot number and port number.

There are 10 slots that contain port blades:

- Slot 5 and 6 contain CP blades.
- Slot 1 through 4 and 7 through 10 contain port blades.

On each port blade, there are 16 or 32 ports (counted from the bottom, 0 to 15, or 0 to 31). A particular port must be represented by both slot number (1 through 4 and 7 through 10) and port number (0 through 15).

When you have port blades with different port counts in the same director (for example, 16-port blade and 32-port blades), the area IDs no longer match the port numbers. Following are the port numbering schemes for the 4/256 SAN Director:

- For the FC4-16 port blade, ports are numbered from 0 through 15 from bottom to top.
- For the FC-32 port blade, ports are numbered from 0 through 15 from bottom to top on the left set of ports and 16 through 31 from bottom to top on the right set of ports.

The Core Switch 2/64 is divided into two logical switches, where slots 1 through 4 constitute logical switch 0 (sw0) and slots 7 through 10 constitute logical switch 1 (sw1). You must be connected to the logical switch that represents the slot where you want to execute a command.

In the SAN Director 2/128 and 4/256 SAN Director default configuration, all the ports are part of a single logical switch. With Fabric OS 4.4.0 and later, you can configure the SAN Director 2/128 as two logical switches (domains).

The following sections tell how to identify ports on the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, and how to identify ports for zoning commands.

By slot and port number

The port number is assigned to an external port to give it a unique identifier in a switch.

To select a specific port in the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, you must identify both the slot number and the port number using the format slot number/port number. No spaces are allowed between the slot number, the slash (/), and the port number.

The following example shows how to enable port 4 on a blade in slot 2:

```
switch:admin> portenable 2/4
```

By port area ID

Zoning commands require that you specify ports using the area ID method. In Fabric OS 4.0.0 and later, each port on a particular domain is given a unique area ID. The relationship between the port number and area ID depends upon the PID format used in the fabric:

- When Core PID format is in effect, the area ID for port 0 is 0, for port 1, it is 1, and so forth.
When using Core PID mode on the Core Switch 2/64 (two logical 64-port switches), 4/256 SAN Director (one domain only), and the SAN Director 2/128 configured with two domains, the area IDs for both logical switches (domains) range from 0 to 63. This means that both logical switch 0 and logical switch 1 have a port that is referenced with area ID 0.
For 32-port blades in the 4/256 SAN Director (using FC4-32), the numbering is contiguous up to port 15; from port 16, the numbering is still contiguous, but you must add 128 to each port number. For example, port 16 in slot 1 has a port number and area ID of 128; port number 15 has a port number and area ID of 15.
- When Extended Edge PID format is in effect, the area ID is the port number plus 16 for ports 0 to 111. For port numbers greater than 111, the area ID wraps around so that port 112 has an area ID of 0, and so on. Each 64-port logical switch (domain) has area IDs ranging from 16 to 79. Port numbers greater than 128 are mapped directly to the core PID.
For details about port area IDs in Extended Edge PID mode, see ["Changing to Extended Edge PID format"](#) on page 220.
- If you perform a port swap operation, the port number and area ID no longer match.

To determine the area ID of a particular port, issue the `switchShow` command. This command displays all ports on the current (logical) switch and their corresponding area IDs.

Basic blade management

This section provides procedures for powering a port blade off and on and for disabling and enabling a port blade.

Powering port blades off and on

Port blades are powered on by default.

Powering off a port blade

1. Connect to the switch and log in as admin.
2. Issue the `slotPowerOff` command with the slot number of the port blade you want to power off.
The slot must exist in the logical switch where you are logged in. For example:

```
switch:admin> slotpoweroff 3
Slot 3 is being powered off
switch:admin>
```

Providing power to a port blade

1. Connect to the switch and log in as admin.
2. Issue the `slotPowerOn` command with the slot number of the port blade you want to power on.
The slot must exist in the logical switch where you are logged in. For example:

```
switch:admin> slotpoweron 3
Powering on slot 3
switch:admin>
```

Disabling and enabling port blades

Port blades are enabled by default.

You might need to disable a port blade to perform diagnostics. When diagnostics are executed manually (from the Fabric OS command line), many commands require the port blade to be disabled. This ensures that diagnostic activity does not interfere with normal fabric traffic.

Disabling a port blade

1. Connect to the switch and log in as admin.
2. Issue the `slotOff` command with the slot number of the port blade you want to disable.

For example:

```
switch:admin> slotoff 3
Slot 3 is being disabled
switch:admin>
```

Enabling a port blade

1. Connect to the switch and log in as admin.
2. Issue the `slotOn` command with the slot number of the port blade you want to enable.


For example:

```
switch:admin> sloton 3
Slot 3 is being enabled
switch:admin>
```

Conserving power

To conserve power and ensure that more critical components are the least affected by a power fluctuation, you can power off components in a specified order, using the `powerOffListSet` command.

The available power is compared to the power demand to determine whether there is enough power to operate. If there is less power available than the demand, the power-off list is processed until there is enough power for operation. By default, the processing proceeds from slot 1 to the last slot in the chassis. As power becomes available, slots are powered up in the reverse order.

 **NOTE:** Some FRUs in the chassis may use significant power, yet they cannot be powered off through software. For example, a missing blower FRU may change the power computation enough to affect how many slots can be powered up.

The `powerOffListShow` command displays the power-off order.

Blade terminology and compatibility

Before configuring a chassis, familiarize yourself with the director CP blade and port blade nomenclature, as well as the port blade compatibilities. Often in procedures, only the abbreviated names for CP and port blades are used (for example, the FC4-16 blade). [Table 15](#) provides CP and port blade abbreviations and descriptions.

Table 15 HP StorageWorks director terminology and abbreviations

Term	Abbreviation	Blade ID	Definition
Core Switch 2/64 CP blade	CP1	1	The first-generation CP blade provided with the Core Switch 2/64. This CP supports 1- and 2-Gbit/sec port speeds. It supports only the dual domain configuration within the chassis.
SAN Director 2/128 CP blade	CP2	5	The second-generation CP blade provided with the SAN Director 2/128. This CP supports 1- and 2-Gbit/sec port speeds. It supports both the dual domain and a single domain configuration within the chassis.
4/256 SAN Director CP blade	CP4	16	The second-generation CP blade provided with the 4/256 SAN Director. This CP supports 1-, 2-, and 4-Gbit/sec port speeds, as well as 16 and 32-port blades.
16-port 2-Gbit/sec port blade	FC-16	2	The first-generation HP StorageWorks director 16-port blade supporting 2-Gbit/sec port speeds. This port blade is compatible only with the Core Switch 2/64 or SAN Director 2/128 CP blades.
16-port 2-Gbit/sec port blade	FC2-16	4	The second-generation HP StorageWorks director 16 port blade supporting 1- and 2-Gbit/sec port speeds. This port blade is compatible only with the SAN Director 2/128 or 4/256 SAN Director CP blades.
16-port 4-Gbit/sec port blade	FC4-16	17	The third-generation HP StorageWorks director 16 port blade supporting 1-, 2-, and 4-Gbit/sec port speeds. This port blade is compatible only with the SAN Director 2/128 and 4/256 SAN Director CP blades. FC4-16 blades do not support private devices.
32-port 4-Gbit/sec port blade	FC4-32	18	A 32-port HP StorageWorks director port blade supporting 1-, 2-, and 4-Gbit/sec port speeds. This port blade is compatible only with the 4/256 SAN Director CP blades. FC4-32 blades do not support private devices.

CP blades

CP blades determine the director type:

- If CP1 blades are installed, the director is a Core Switch 2/64.
- If CP2 blades are installed, the director is a SAN Director 2/128.
- If CP4 blades are installed, the director is a 4/256 SAN Director.

HP recommends that each HP StorageWorks director have only one type of CP blade installed and that each CP (primary and secondary partition) maintains the same firmware version.

Port blade compatibility

Table 16 indicates which blades are supported for each HP StorageWorks director.

Table 16 Blades supported by each HP StorageWorks director

Director	Port Blades			
	FC-16	FC2-16	FC4-16	FC4-32
Core Switch 2/64 (CP1)	Supported	N/A	N/A	N/A
SAN Director 2/128 (CP2)	Supported	Supported	Supported	N/A
4/256 SAN Director (CP4)	N/A	Supported	Supported	Supported

Setting chassis configurations

The `chassisConfig` command allows you to set the chassis configuration for products that support both single-switch (one domain) and dual-switch (two domains) operation.

Table 17 lists the supported configurations options for Fabric OS 5.x.

Table 17 Supported configuration options

Option	Number of domains	Maximum number of ports per switch	Supported port blades	Supported CP blades	Notes	Results
1	1	128	FC2-16, FC4-16	CP2 or CP4	CP4 fits all chassis except the D2 chassis. Option 1 is the default configuration for the SAN Director 2/128.	One 128-port switch (Blade IDs 4, 17 on slots 1–4, 7–10. Blade ID 5 and 16 on slots 5, 6)
2	2	64/64	FC2-16	CP2	N/A	Two 64-port switches (Blade ID 4 on slots 1–4, 7–10. Blade ID 5 on slots 5, 6)
3	2	64/64	Left side: FC2-16 Right side: FC-16	CP2	N/A	Two 64-port switches (Blade ID 4 on slots 1–4; ID 2 on slots 7–10. Blade ID 5 on slots 5, 6)

Table 17 Supported configuration options (continued)

Option	Number of domains	Maximum number of ports per switch	Supported port blades	Supported CP blades	Notes	Results
4	2	64/64	Left side: FC-16 Right side: FC2-16	CP2	N/A	Two 64-port switches (Blade ID 2 on slots 1–4; ID 4 on slots 7–10. Blade ID 5 on slots 5, 6)
5	1	256	FC4-16, FC4-32	CP4	CP4 fits all chassis except the D2 chassis. Option 5 is the default configuration option for 4/256 SAN Director.	One 256-port switch (Blade IDs 4, 17, and 18 on slots 1–4, 7–10. Blade ID 16 on slots 5, 6)

The following sections contain procedures for obtaining chassis information and for configuring director domains using the `chassisConfig` command.

Obtaining slot information

For a Core Switch 2/64 or SAN Director 2/128 configured as two logical switches, the chassis-wide commands display or control both logical switches. In the default configuration, the SAN Director 2/128 and 4/256 SAN Director are configured as one logical switch, so the chassis-wide commands display and control the single logical switch.

Displaying the status of all slots in the chassis

1. Connect to the switch and log in as user or admin.
2. Issue the `slotShow` command to display the current status of each slot in the system.

The format of the display includes a header and four fields for each slot. The fields and their possible values are:

Field	Value
Slot	Displays the physical slot number.
Blade Type	Displays the blade type: <ul style="list-style-type: none"> • SW BLADE: The blade is a switch. • CP BLADE: The blade is a CP. • UNKNOWN: The blade is not present or its type is not recognized.
ID	Displays the hardware ID of the blade type.

Field	Value
Status	<p>Displays the status of the blade:</p> <ul style="list-style-type: none"> • VACANT: The slot is empty. • INSERTED, NOT POWERED ON: The blade is present in the slot but is turned off. • DIAG RUNNING POST1: The blade is present, powered on, and running the post-initialization POST. • DIAG RUNNING POST2: The blade is present, powered on, and running the POST. • ENABLED: The blade is on and enabled. • ENABLED (User Ports Disabled): The blade is on, but external ports have been disabled with the <code>bladeDisable</code> command. • DISABLED: The blade is powered on but disabled. • FAULTY: The blade is faulty because an error was detected. The reason code numbers displayed are for debugging purposes. • UNKNOWN: The blade is inserted but its state cannot be determined.

Configuring a new SAN Director 2/128 with two domains

By factory default, the SAN Director 2/128 is configured as one 128-port switch (one domain). The following procedure assumes that the new director:

- Has been installed and connected to power, but is not yet attached to the fabric.
- Has been given an IP address, but is otherwise running factory defaults. If this is not the case, back up the current configuration before starting, so that you can restore it later if necessary.
- Is running Fabric OS 4.4.0 or later.
- Is running in configuration option one (one switch, FC2-16 cards installed).

Use the following procedure to add a factory-new SAN Director 2/128 to a fabric and configure it as two 64-port switches (two domains).

1. Connect to the switch and log in as admin.
2. Issue the `chassisconfig` command without options to verify that the switch is configured with one domain.

For example:

```
chassisconfig
Current Option: 1
```

3. Issue the `chassisconfig` command to configure two domains. Use the `-f` option to suppress prompting for uploading the configuration.

This command reboots the system. For example:


```
chassisconfig -f 2
Current Option changed to 2
Restoring switch 0 configuration to factory defaults...
All account passwords have been successfully set to factory default.
Restoring switch 1 configuration to factory defaults...
All account passwords have been successfully set to factory default.
```

4. After the system reboots, log in again to the first logical switch (sw0) as admin.

5. Use the `configure` command to configure the sw0 to match your fabric specifications.
If the director is to be merged into an existing fabric, do not configure zoning parameters; these are propagated when you merge the director into the fabric.
6. Log in to the second logical switch (sw1) as admin.
7. Use the `configure` command to configure the sw1 to match your fabric specifications.
If the director is to be merged into an existing fabric, do not configure zoning parameters; these are propagated when you merge the director into the fabric.
8. If the fabric is in secure mode, perform the following steps; otherwise, proceed to [step 9](#).
(See the *HP StorageWorks Secure Fabric OS administrator guide* for specific instructions.)
 - a. Optional: To configure sw0 and sw1 in one operation, connect them with an ISL link to form a temporary fabric.
 - b. If you want sw0 and sw1 to be FCSs, update the overall fabric's FCS policy to include them. If not, skip this step.
 - c. On sw0, enable security mode and use the `secModeEnable` command to create an FCS list that matches your overall fabric's FCS policy.
 - d. Reset the version stamp on sw0.
 - e. If you connected sw0 and sw1 in [step 8a](#) and you do not want them connected, disconnect the ISL link between them. If you did not connect them, repeat [step 8b](#) through [step 8d](#) on sw1.
9. Optional: Connect the new two-domain SAN Director 2/128 to the fabric.
10. Issue the `fabricShow` command to verify that sw0 and sw1 have been merged with the fabric.
11. Issue the `cfgShow` command to verify that zoning parameters were propagated.

Converting an installed SAN Director 2/128 to support two domains

Fabric OS versions earlier than 4.4.0 supported only one domain for SAN Director 2/128s (one 128-port logical switch). When you upgrade a SAN Director 2/128 to Fabric OS 4.4.0 or later, you can use the `chassisConfig` command to specify two domains for the director (two 64-port logical switches, sw0 and sw1). This conversion is for SAN Director 2/128s using configuration option one (one switch, FC2-16 cards installed).

 **NOTE:** This procedure restores most configuration parameters to factory defaults. After performing this procedure, you must check the new configuration and reconfigure those parameters that you customized in the old configuration.

During this procedure, power is reset and the CP blades are rebooted, so traffic on the fabric is disrupted. If the fabric is in secure mode, enabling security on the new domains is a complicated task. Do not convert existing core switches.

1. Connect to the switch and log in as admin.
2. If the director is already in a fabric, minimize disruption by removing the director from the fabric using one of the following methods:
 - Physically disconnect the director.
 - Use the `portCfgPersistentDisable` command on all connected remote switches to persistently disable their ports that are connected to the director, or remove ISLs that connect the SAN Director 2/128 to the current fabric.
3. Issue the `chassisConfig` command to change the configuration from the default (one domain) to two domains. The following command reboots the system:
`chassisconfig 2`

During the conversion, you are prompted to save the configuration of sw0. Follow the prompts to save the configuration file.

4. Issue the `ipAddrSet` command to set and confirm the IP address of sw1 (sw1 takes on a default that must be corrected).
The IP address of sw0 is already set.
5. After the system reboots, log in again as admin to each logical switch and issue the `switchName` command to assign a name to the new switch.
6. Using the configuration file saved in [step 3](#) as a guide, manually reconfigure sw0 and sw1.
Do not configure zoning parameters; these are propagated when you merge the director into the fabric.
7. If the fabric is in secure mode, perform the following steps; otherwise, proceed to [step 8](#).
 - a. Optional: to configure sw0 and sw1 in one operation, connect them with an ISL link to form a temporary fabric.
 - b. If you want sw0 and sw1 to be FCSs, update the overall fabric's FCS policy to include them. If not, skip this step.
 - c. On sw0, enable security mode and use the `secModeEnable` command to create an FCS list that matches your overall fabric's FCS policy.
 - d. Reset the version stamp on sw0.
 - e. If you connected sw0 and sw1 in [step 7a](#) and you do not want them connected, disconnect the ISL link between them. If you did not connect them, repeat [step 7b](#) through [step 7d](#) on sw1.
8. If you physically disconnected the switch in [step 2](#), reconnect it to the fabric.
If you used the `portCfgPersistentDisable` command in [step 2](#), use the `portCfgPersistentEnable` command to persistently enable all ports that connect the switch to other switches in the fabric.
9. Issue the `fabricShow` command to verify that sw0 and sw1 have been merged with the fabric.
10. Issue the `configShow` command to verify that zoning parameters were propagated.

Setting the blade beacon mode

When beaconing mode is enabled, the port LEDs flash amber in a running pattern from port 0 through port 15 and back again. The pattern continues until you turn beaconing mode off. Use the flashing LEDs to locate a particular blade.

Setting the blade beacon mode on:

1. Connect to the switch and log in as admin.
2. Issue the `bladeBeacon` command:

```
bladebeacon slotnumber, mode
```

The *slotnumber* is the blade on which you want to enable beacon mode; this slot number must exist on the logical switch. A *mode* value of 1 turns beaconing on, and 0 turns beaconing off. For example:

```
switch:admin> bladebeacon 3, 1
switch:admin>
```


6 Routing traffic

This chapter describes HP StorageWorks switch routing features and procedures.

About data routing and routing policies

Data moves through a fabric from switch to switch and from storage to server along one or more paths that make up a route. Routing policies determine the correct path for each frame of data.

- △ **CAUTION:** For most configurations, the default routing policy is optimal, and provides the best performance. Change the policy only if there is a performance issue that is of concern, or a particular fabric configuration requires it.

The following routing policies are available to tune routing performance:

- **Exchange-based routing:** The choice of routing path is based on the source ID (SID), destination ID (DID), and Fibre Channel originator exchange ID (OXID), optimizing path utilization for the best performance. Thus, every exchange can take a different path through the fabric
- **Device-based routing:** The choice of routing path is based on the Fibre Channel addresses of the SID and the DID, improving path utilization for better performance. Thus, the same route is always used and the sequence of exchanges is guaranteed.
- **Port-based routing:** The choice of routing path is based only on the incoming port and the destination domain. To optimize port-based routing, the Dynamic Load Sharing feature (DLS) can be enabled to balance the load across the available output ports within a domain.

Device-based and exchange-based routing require the use of DLS; when these policies are in effect, you cannot disable the DLS feature.

Using port-based routing, you can assign a *static route*, in which the path chosen for traffic never changes. In contrast, device-based and exchange-based routing policies always employ *dynamic path selection*. Port-based routing is supported by all HP StorageWorks models (except the 4/256 SAN Director using configuration option 5; see [Table 17](#) on page 91).

Specifying the routing policy

The following routing policies are supported:

- **Port-based path selection:** The default on SAN Switch 2/8V, SAN Switch 2/16V, and SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director (using configuration options 1 through 4). These switches support the port-based policy only; you cannot change the routing policy for these switches. The 4/8 SAN Switch, 4/16 SAN Switch, 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32 can also use port-based routing.
- **Device-based path selection:** Available on 4/8 SAN Switch, 4/16 SAN Switch, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32, and 4/256 SAN Director (using configuration option 5). If there are devices in your fabric that cannot accommodate out-of-order exchanges, use the device-based policy. In FICON environments device-based routing is recommended.
- **Exchange-based path selection:** The default on the 4/8 SAN Switch, 4/16 SAN Switch, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32, and 4/256 SAN Director (using configuration option 5).

See ["Configuring Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director"](#) on page 87 for details about 4/256 SAN Director configuration options.

You can use the `aptPolicy` command to display and specify a different routing policy. Note that if you attempt to set the policy when the 4/256 SAN Director uses configuration options 1–4, an error message is returned. See the *HP StorageWorks Fabric OS 5.x command reference guide* for details on the `aptPolicy` command.

You must disable the switch before changing the routing policy, and reenable it afterward.

In the following example, the routing policy for a SAN Switch 4/32 is changed from exchange-based to device-based:


```
switch:admin> aptpolicy
Current Policy: 3

3: Default Policy
1: Port Based Routing Policy
2: Device Based Routing Policy
3: Exchange Based Routing Policy
switch:admin> switchdisable
switch:admin> aptpolicy 2
Policy updated successfully.
switch:admin> switchenable
switch:admin> aptpolicy
Current Policy: 2
```

Assigning a static route

A static route can be assigned only when the active routing policy is port-based. When device-based or exchange-based routing is active, you cannot assign static routes. Thus, the 4/256 SAN Director using configuration option 5 does not support static routing.

To assign a static route, use the `uRouteConfig` command. To remove a static route, use the `uRouteRemove` command.

 **NOTE:** For the SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director (using configuration options 1 through 4):

When you issue the `uRouteConfig` command, two similar warning messages might be displayed if a platform conflict occurs. The first message is displayed when the static routing feature detects the conflict. The second message is displayed when the DLS feature detects the condition as it tries to rebalance the route.

A platform conflict occurs if a static route was configured with a destination port that is currently down. The static route is ignored in this case, in favor of a normal dynamic route. When the configured destination port comes back up, the system attempts to reestablish the static route, potentially causing a conflict.

Specifying frame order delivery

The order of delivery of frames is maintained within a switch and determined by the routing policy in effect. Following are the frame delivery behaviors for each routing policy.

- **Port-based routing:** All frames received on an ingress port destined for a destination domain are guaranteed to exit the switch in the same order in which they were received.
- **Device-based routing:** All frames received on an ingress port between the same two fabric devices are guaranteed to exit the switch in the same order in which they were received. This policy maintains the order of frames across exchanges between the fabric devices as well.
- **Exchange-based routing:** All frames received on an ingress port for a given exchange are guaranteed to exit the switch in the same order in which they were received. Because different paths are chosen for different exchanges, this policy does not maintain the order of frames across exchanges.


If even one switch in the fabric delivers out-of-order exchanges, exchanges are then delivered to the target out-of-order, regardless of the policy configured on other switches in the fabric.

In a stable fabric, frames are always delivered in order, even when the traffic between switches is shared among multiple paths. However, when topology changes occur in the fabric (for example, if a link goes down), traffic is rerouted around the failure, and some frames could be delivered out of order. Most destination devices tolerate out-of-order delivery, but some do not.

By default, out-of-order frame-based delivery is allowed to minimize the number of frames dropped. Force in-order frame delivery only across topology changes if the fabric contains destination devices that cannot tolerate occasional out-of-order frame delivery.

Forcing in-order frame delivery across topology changes

1. Connect to the switch and log in as admin.
2. Issue the `iodSet` command.

 **NOTE:** This command can cause a delay in the establishment of a new path when a topology change occurs; use it with care.

Restoring out-of-order frame delivery across topology changes

1. Connect to the switch and log in as admin.
2. Issue the `iodReset` command.

Using DLS

The device-based and exchange-based routing policies depend on the Fabric OS DLS feature for dynamic routing path selection. When these policies are in force, DLS is by default enabled and cannot be disabled.

When the port-based policy is in force, you can enable DLS to optimize routing. When DLS is enabled, it shares traffic among multiple equivalent paths between switches. DLS recomputes load sharing when a switch boots up, each time an E_Port goes offline and online, or when an Fx_Port goes offline.

Checking and setting DLS

1. Connect to the switch and log in as admin.
2. Issue the `dlsShow` command to view the current DLS setting.
One of the following messages appears:
 - `DLS is set`, which indicates that DLS is turned on.
 - `DLS is not set`, which indicates that DLS is turned off.
3. Issue the `dlsSet` command to enable DLS or issue the `dlsReset` command to disable it.
You cannot disable DLS when device-based or exchange-based routing policies are in effect.
For example:

```
switch:admin> dlsShow
DLS is not set
switch:admin> dlsSet
switch:admin> dlsShow
DLS is set
switch:admin> dlsreset
switch:admin> dlsShow
DLS is not set
```

Viewing routing path information

The `topologyShow` and `uRouteShow` commands provide information about the routing path.

1. Connect to the switch and log in as admin.
2. Issue the `topologyShow` command to display the fabric topology, as it appears to the local switch.

The following entries appear:

Local Domain ID	Domain number of the local switch
Domain	Domain number of the destination switch
Metric	Cost of reaching the destination domain
Name	The name of the destination switch
Path Count	The number of currently active paths to the destination domain
Hops	The maximum number of hops to reach the destination domain.
Out Port	The Port to which the incoming frame will be forwarded in order to reach the destination domain
In Ports	Input ports that use the corresponding Out Port to reach the destination domain
Total Bandwidth	The maximum bandwidth of the out port
Bandwidth Demand	The maximum bandwidth demand of the in ports
Flags	Always D, indicating a dynamic path

For example:

```
switch:admin> topologyshow
2 domains in the fabric; Local Domain ID: 1
Domain: 6
Metric: 500
Name: switch
Path Count: 4
Hops: 1
Out Port: 60
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 61
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 62
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
Hops: 1
Out Port: 58
In Ports: None
Total Bandwidth: 2 Gbps
Bandwidth Demand: 0%
Flags: D
```

3. Issue the `uRouteShow` command to display unicast routing information.

For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32, use the following syntax:

```
urouteshow [portnumber][, domainnumber]
```

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director: Use the following syntax:

```
urouteshow [slot/][portnumber][, domainnumber]
```

The following entries appear:

- **Local Domain:** The domain number of the local switch.
- **In Ports:** Port from which a frame is received.
- **Domain:** The destination domain of the incoming frame.
- **Out Port:** The port to which the incoming frame is forwarded in order to reach the destination domain.
- **Metric:** The cost of reaching the destination domain.
- **Hops:** The maximum number of hops to reach the destination domain.
- **Flags:** Indicates whether the route is dynamic (D) or static (S). A static route is assigned using the `uRouteConfig` command.
- **Next (Dom, Port):** The domain number and port number of the next hop.

The following example displays the routing information of all the active ports:

```
switch:admin> urouteshow
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
0      1  11 1000 1 D  1,0
11     2   0 1500 2 D  4,0
      4  16 500 1 D  4,0
16     1  27 1000 1 D  1,1
27     2  16 1500 2 D  4,16
4      0  29 500 1 D  4,0
```

The following example displays the routing information for port 11 on slot 1:

```
switch:admin> urouteshow 1/11
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
11     2   0 1500 2 D  4,0
      4  16 500 1 D  4,0
```

The following example displays the routing information of port 11 to domain 4 only:

```
switch:admin> urouteshow 1/11, 4
Local Domain ID: 3
In PortDomain Out Port Metric Hops Flags Next (Dom, Port)
-----
11     4  16 500 1 D  4,0
```

Viewing routing information along a path

You can display detailed routing information from a source port (or area) on the local switch to a destination port (or area) on another switch. This routing information describes the full path that a data stream travels between these ports, including all intermediate switches.

1. Connect to the switch and log in as admin.
2. Issue the `pathInfo` command.

In interactive mode, you can specify the following parameters for display:

- `Max hops`: The maximum number of hops that the `pathinfo` frame is allowed to traverse.
- `Domain`: The destination domain ID.
- `Source Port`: The port number (or area number for Core Switch 2/64, SAN Director 2/128, or 4/256 SAN Director) on which the switch receives frames.
- `Destination Port`: The output port that the frames use to reach the next hop on this path. For the last hop, the destination port.
- `Basic stats`: Basic statistics on every link.
- `Extended stats`: Detailed statistics on every link.
- `Trace reverse path`: Traverses from the destination switch back to the source switches.
- `Source route`: Forces the frame to follow a specified path to reach the destination.
- `Timeout`: The maximum time to wait for a response from `pathInfo`, in seconds.

Paths always originate on the local switch. The path destination can be specified by domain or port. By default, the path is the path taken by traffic from the source to destination port, but you can also specify all or portions of a path.

See the *HP StorageWorks Fabric OS 5.x command reference guide* for details on the `pathInfo` command.

The following example is from a SAN Switch 2/32 (other models provide similar information):

```
switch:admin> pathinfo

Max hops: (1..127) [25]
Domain: (1..239) [-1] 1
Source port: (0..255) [-1]
Destination port: (0..255) [-1]
Basic stats (yes, y, no, n): [no]
Extended stats (yes, y, no, n): [no]
Trace reverse path (yes, y, no, n): [no]
Source route (yes, y, no, n): [no]
Timeout: (1..30) [10]

Target port is Embedded

Hop  In Port  Domain ID (Name)          Out Port  BW    Cost
-----
0      E        10 (SW3900)              15        2G    500
1      7         1 (swd3900TechPu        E         -      -
switch:admin>
```

The information that `pathInfo` provides is as follows:

- `Hop`: The hop number. The local switch is hop 0.
- `In Port`: The port that the frames come in from on this path. For hop 0, the source port.
- `Domain ID`: The domain ID of the switch.

- `Name`: The name of the switch.
- `Out Port`: The output port that the frames use to reach the next hop on this path. For the last hop, the destination port.
- `BW`: The bandwidth of the output ISL, in Gbit/sec. It does not apply to the embedded port.
- `Cost`: The cost of the ISL used by FSPF routing protocol. It applies only to an E_Port.

7 Administering FICON fabrics

 **NOTE:** FICON is not supported on HP B-Series Fibre Channel switches. The FICON information in this document is included for reference only.

FICON overview

IBM FICON is an industry-standard, high-speed input/output (I/O) interface for mainframe connections to storage devices. Fabric OS supports *intermix mode* operations, in which FICON and Fibre Channel technology work together. For specific information about intermix mode and other aspects of FICON, see the IBM Redbook, *FICON Native Implementation and Reference Guide*.

Fabric OS provides standard support for FICON single-switch operation.

Multiple-switch cascaded FICON operation requires an HP Secure Fabric OS license.

CUP operation requires an HP FICON CUP license.

The following Fabric OS standard features support FICON fabrics:

- **Port swapping:** Redirects resources from a failed port to a healthy port without changing the FICON host configuration. Port swapping is available for both FICON and open system environments. Port swapping resolves situations in which the hardware has failed and the channel configurations cannot be changed quickly. Port swapping has minimal or no impact on other switch features.
- **Insistent domain ID (IDID):** Allows the switch to insist on a specific domain ID before joining a fabric. This feature guarantees that a switch operates only with its preassigned domain ID.
- **The FICON MIB module:** Addresses link incident data for FICON hosts and devices connected to a switch. It supplements other MIBs used to manage switches and should be used with those other MIBs. For more information, see the *HP StorageWorks Fabric OS 5.x MIB reference guide*.
- **Link incident detection, registration, and reporting:** Provides administrative and diagnostic information.

The following optional features provide further support:

- **The Secure Fabric OS optional license:** Includes fabric binding, switch binding, and port binding security methods that prevent unauthorized devices from joining a fabric.
- **A Fabric Manager optional license:** Can be used to manage a fabric that supports FICON and Fibre Channel Protocol (FCP) devices and traffic. This is the recommended GUI management tool for FICON environments.
- **Advanced Web Tools:** Can be used to manage a director (switch) that supports FICON and FCP devices and traffic.

To incorporate and manage FICON on a switch or fabric, your system must have Fabric OS 4.1.2 or later installed. If you are implementing FICON in a single-switch non-candidate environment, there are no additional software requirements. The Secure Fabric OS and Advanced Zoning optional licensed features are required on all switches participating in a FICON multiple-switch cascaded environment.

 **NOTE:** Some licenses are installed and activated on the switch at the factory. Use an HP management interface to verify that the required licenses are installed and activated on the switch.

The optional Secure Fabric OS license provides the following fabric, switch, and port binding features:

- Fabric binding is a security method for restricting switches within a multiple-switch fabric. The Switch Connection Control (SCC) policy prevents unauthorized switches from joining a fabric. Switches are authenticated using digital certificates and unique private keys provided to the Switch Link Authentication Protocol (SLAP).

- Switch binding is a security method for restricting devices that connect to a particular switch. If the device is another switch, security handled by the SCC policy. If the device is a host or storage device, the Device Connection Control (DCC) policy binds those devices to a particular switch. Policies range from completely restrictive to reasonably flexible, based upon customer needs.
- Port binding is a security method for restricting host or storage devices that connect to particular switch ports. The DCC policy also binds device ports to switch ports. Policies range from completely restrictive to reasonably flexible, based upon customer needs.

There are two types of FICON configurations:

- A *single-switch* configuration (called *switched point-to-point*) requires that the channel be configured to use single-byte addressing. If the channel is set up for two-byte addressing, the cascaded configuration setup applies. This type of configuration is described in "[Configuring a single switch](#)" on page 108.
- A *cascaded configuration* (known as a *high-integrity fabric*) requires a list of authorized switches. This authorization feature (called *fabric binding*) is available through Secure Fabric OS. The fabric binding policy allows a predefined list of switches (domains) to exist in the fabric and prevents other switches from joining the fabric. This type of configuration is described in "[Configuring a high-integrity fabric](#)" on page 109.

CUP protocol is used by IBM mainframe management programs to provide in-band management for FICON switches. When it is enabled, you can set up directors in a FICON environment to be managed through IBM mainframe management programs. CUP is an optionally licensed feature available with Fabric OS 4.4.0 or later.

[Table 18](#) summarizes the Fabric OS CLI commands that can be used for managing FICON fabrics. For detailed information on these commands, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

Table 18 Fabric OS commands related to FICON and FICON CUP

Command	Description
Standard Fabric OS commands:	
configure	Sets the domain ID and the IDID mode.
portSwap	Swaps ports.
portSwapDisable	Disables the portSwap command.
portSwapEnable	Enables the portSwap command.
portSwapShow	Displays information about swapped ports.
Commands specific to FICON:	
ficonclear rlir	Removes all registered link incident records (RLIRs) from the local RLIR database.
ficonclear rnid	Removes all outdated RNID records from the local RNID database.
ficonshow ilir [fabric]	Displays FRU failure information on the local switch or on the fabric.
ficonshow lirr [fabric]	Displays registered listeners for link incidents for the local switch or for the fabric.
ficonshow rlir [fabric]	Displays link incidents for the local switch or for the fabric.
ficonshow rnid [fabric]	Displays node identification data for all devices registered with the local switch or all devices registered with all switches defined in the fabric.
ficonshow switchrnid [fabric]	Displays node identification data for the local switch or for the fabric.
Commands specific to FICON CUP:	
ficoncupset fmsmode	Sets FICON Management Server mode on or off for the switch.
ficoncupset modereg	Sets the mode register bits for the switch.
ficoncupshow fmsmode	Displays the FICON Management Server mode setting for the switch.
ficoncupshow modereg	Displays the mode register bit settings for the switch.

The Fabric OS CLI supports only a subset of the HP management features for FICON fabrics. The full set of FICON CUP administrative procedures is available using the HP Fabric Manager and Advanced Web Tools software features. You can also use an SNMP agent and the FICON MIB. For information on these tools, see:

- Advanced Web Tools: *HP StorageWorks Fabric OS 5.x Advanced Web Tools administrator guide*
- Fabric Manager: *HP StorageWorks Fabric Manager 5.x administrator guide*
- SNMP Agent and FICON MIB: *HP StorageWorks Fabric OS 5.x MIB administrator guide*

Configuring switches

This section describes how to configure a switch in a FICON environment. Use the worksheet shown in [Figure 3](#) on page 119 to record your configuration information.

The following are recommended FICON environment configuration settings:

- Disable DLS (`dlsReset` command).

If DLS is enabled, traffic on existing ISL ports might be affected when one or more new ISLs are added between the same two switches. Specifically, adding a new ISL might result in dropped frames as routes are adjusted to take advantage of the bandwidth provided. By disabling DLS, you ensure that there are no dropped frames.

A similar situation occurs when an ISL port is taken offline and then brought back online. When the ISL port goes offline, the traffic on that port is rerouted to another ISL with a common destination. When the ISL port comes back online and DLS is enabled, the rerouting of traffic back to the ISL port might result in dropped frames. If DLS is not enabled, traffic is not routed back.

- Configure ports that are connected to 1-Gbit/sec channels for fixed 1-Gbit/sec speed. Otherwise, when using fixed 1-Gbit/sec channels (both G5 and FICON Express), the FICON host might generate erroneous link incidents when the channels are coming online. These link incidents result in a Call Home. Other than the generated link incident, the channel comes online and functions normally.
- Enable in-order delivery (`iodSet` command).
- Enable VC translation link initialization on extended fabrics links, to stabilize them. See ["For dynamic long-distance links, you can approximate the number of buffer credits using the following formula:"](#) on page 165 for details on this option of the `portCfgLongDistance` command.
- Although there are no specific zoning rules related to FICON environments, HP recommends that you follow standard FCP zoning practices. For management purposes, put FCP devices in one zone and FICON devices in another zone when operating in a mixed environment.

Preparing a switch

To verify that a switch is ready to be used in a FICON environment, complete the following steps:

1. Connect to the switch and log in as admin.
2. If the switch is not in a cascaded environment, proceed to [step 3](#).

If the switch is in a FICON cascaded environment, issue the following commands:

- `licenseShow` to verify that required licenses (Secure Fabric OS and Zoning) are activated.
- `secModeShow` to determine whether Secure Fabric OS is enabled; if it is disabled, enable it.
- `secPolicyShow` to verify that the SCC_POLICY is active.

3. Issue `switchShow` to verify that the switch and devices are online.
4. Issue `ficonshow rnid` to verify that the FICON devices are registered with the switch.
5. Issue `ficonshow lirr` to verify that the FICON host channels are registered to listen for link incidents.
6. Optional: To use FICON CUP, see ["Using FICON CUP"](#) on page 111.

Configuring a single switch

Single-switch configuration does not require IDID or fabric binding, provided that connected channels are configured for single-byte addressing. However, configure IDID to ensure that domain IDs are maintained.

Configuring a high-integrity fabric

To configure a high-integrity fabric (cascaded configuration):

1. Disable each switch in the fabric.
2. For each switch:
 - a. Enable the IDID flag.
 - b. Set the domain ID.
 - c. Install security certificates and keys.
3. Enable the switches.
This builds the fabric.
4. Set up security on the primary FCS switch.
Use Quickmode for FICON, which activates the SCC_POLICY and does not create a DCC policy. The security policies are distributed to each switch in the fabric. (For details on the Quickmode procedure, see the *HP StorageWorks Secure Fabric OS administrator guide*.)
5. Connect and enable channel and control unit (CU) devices.
The Query for Security Attributes (QSA) response to the channel indicates that the fabric binding and IDID are enabled.

Figure 1 and Figure 2 show two viable cascaded configurations. These configurations require Channel A to be configured for two-byte addressing and require IDID and fabric binding. There can be only two switches in the path from the channel to the CU.

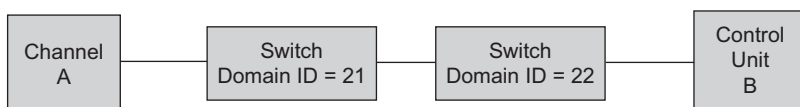


Figure 1 Cascaded configuration with two switches

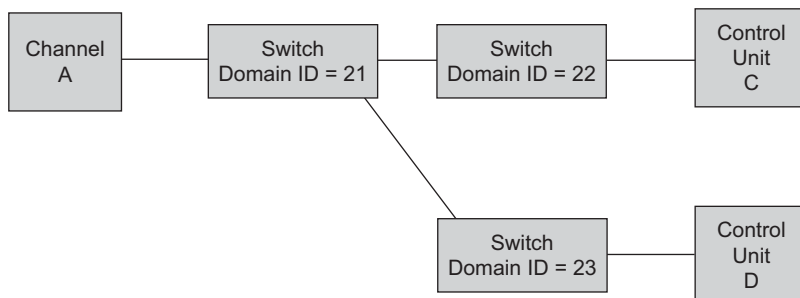


Figure 2 Cascaded configuration with three switches

Setting a unique domain ID

In a cascaded configuration, each switch must have a unique domain ID, and insistent IDID mode must be enabled. To set a unique domain ID and enable IDID mode, follow these steps:

1. Connect to the switch and log in as admin.
2. Verify that the switch has a unique domain ID. If it does not, set a unique domain ID.
For instructions on displaying and changing the domain ID, see "[Working with domain IDs](#)" on page 32.
3. Issue the `switchDisable` command to disable the switch.
4. Issue the `configure` command.
5. Enter `y` after the Fabric Parameters prompt.
6. To enable IDID mode, enter `y` after the Insistent Domain ID Mode prompt.
(You can disable this mode by entering `n`.)

7. Respond to the remaining prompts (or press **Ctrl-d** to accept the other settings and exit).
8. Issue the `switchEnable` command to reenable the switch.

For example:

```
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [3] 5
    R_A_TOV: (4000..120000) [10000]
    E_D_TOV: (1000..5000) [2000]
    Data field size: (256..2112) [2112]
    Sequence Level Switching: (0..1) [0]
    Disable Device Probing: (0..1) [0]
    Suppress Class F Traffic: (0..1) [0]
    VC Encoded Address Mode: (0..1) [0]
    Per-frame Route Priority: (0..1) [0]
    Long Distance Fabric: (0..1) [0]
    BB credit: (1..16) [16]
Insistent Domain ID Mode (yes, y, no, n): [no] y
Virtual Channel parameters (yes, y, no, n): [no]
Switch Operating Mode (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
Portlog events enable (yes, y, no, n): [no]
Committing configuration...done.
switch:admin>
```

Displaying information

You can display link incidents, registered listeners, node identification data, and FRU failures, as described in the following procedures.

Link incidents

The RLIR Extended Link Service (ELS) contains the link incident information sent to a listener N_Port.

To display link incidents, connect to the switch, log in as user, and issue one of the following commands:

- For the local switch: `ficonshow rlir`
- For all switches defined in the fabric: `ficonshow rlir fabric`

Registered listeners

To display registered listeners for link incidents, connect to the switch, log in as user, and issue one of the following commands:

- For the local switch: `ficonshow lirr`
- For all switches defined in the fabric: `ficonshow lirr fabric`

Node identification data

To display node-identification data, connect to the switch, log in as user, and issue any of the following commands:

- For the local switch: `ficonshow switchrnid`
- For all switches defined in the fabric: `ficonshow switchrnid fabric`
- For all devices registered with the local switch: `ficonshow rnid`
- For all devices registered with all switches defined in the fabric: `ficonshow rnid fabric`

FRU failures

To display FRU failure information, connect to the switch, log in as admin, and issue one of the following commands:

- For the local switch: `ficonshow ilir`
- For all switches defined in the fabric: `ficonshow ilir fabric`

Swapping ports

If a port malfunctions, or if you want to connect to different devices without having to rewire your infrastructure, you can move a port's traffic to another port (swap ports) without changing the I/O Configuration Data Set (IOCDS) on the mainframe computer.

To swap ports, perform the following steps (see the example that follows):

1. Connect to the switch and log in as admin.
2. Issue the `portSwapEnable` command (to enable the command for port swapping).
3. Issue the `portDisable` command to disable the two ports to be swapped.
4. Issue the `portSwap` command to swap the ports.
Any port in the switch can be used as the alternate for any other port within the same switch.
5. Reenable the ports using the `portEnable` command.
6. Issue `portSwapDisable` (to disable the command for port swapping).

In the following example:

- `slot` is the slot number of the port blade for a system with port blades (optional).
- `portA` is the original port number.
- `portB` is the alternate port number.

```
switch:admin> portswapenable
switch:admin> portdisable [slot/] portA [slot/]portB
switch:admin> portswap [slot/] portA [slot/]portB
switch:admin> portenable [slot/] portA [slot/]portB
switch:admin> portswapdisable
```

You can use the `portSwapShow` command to display information about swapped ports in a switch.

You can use the `portSwap` command to disable the portswap feature. You cannot use the `portSwap` command after this feature is disabled. The enabled state of the portswap feature is persistent across reboots and power cycles. Enabling and disabling the portswap feature does not affect previously executed portswap operations.

See the *HP StorageWorks Fabric OS 5.x command reference guide* for additional details about the `portSwap` command.

Clearing the FICON management database

You can clear RLIR and RNID records from the FICON management database as follows:

1. Connect to the switch and log in as admin.
2. To remove all the RLIR records from the local RLIR database, issue `ficonclear rlir`.
3. To remove all the RNID records marked not current from the local RNID database, issue `ficonclear rnid`.

Using FICON CUP

Host-based management programs manage switches using CUP protocol by sending commands to an emulated control device in Fabric OS. An HP switch that supports CUP can be controlled by one or more host-based management programs, as well as by HP tools.

A *mode register* controls the behavior of the switch with respect to CUP itself, and with respect to the behavior of other management interfaces.

FICON Management Server mode (fmsmode) must be enabled on the switch to enable CUP management features. When this mode is enabled, Fabric OS prevents local switch commands from interfering with host-based management commands by initiating serialized access to switch parameters.

If more than one switch is to be used in the FICON CUP fabric, Secure Fabric OS must be installed. See ["Configuring a high-integrity fabric"](#) on page 109 for more information.

If HP Advanced Zoning is in use, see ["Zoning and PDCM considerations"](#) on page 116.

Setup summary

To set up FICON CUP, perform the following actions in the order stated:

1. Install Fabric OS 4.4.0 or later on an HP StorageWorks switch.
2. For the SAN Director 2/128 only, use the `portDisable` command to disable (block) port 126.
Port 126 is not supported in a CUP environment. After fmsmode has been successfully enabled, port 126 remains disabled. It cannot be used either as an F_Port or an E_Port. Because port 126 is not available after enabling fmsmode, first move any fiber connected to port 126 to another free port.
3. Install a CUP license on the switch. See ["Maintaining licensed features"](#) on page 26.
4. Enable FICON management server mode (fmsmode) on the switch. See ["Enabling and disabling FMS mode"](#) on page 112.

After completing the setup, you can configure CUP attributes (FMS parameters). See ["Setting mode register bits"](#) on page 114.

Enabling and disabling FMS mode

To enable fmsmode:

1. Connect to the switch and log in as admin.
2. Issue `ficoncupset fmsmode enable`.


To disable fmsmode:

1. Connect to the switch and log in as admin.
2. Issue `ficoncupset fmsmode disable`.

The fmsmode setting can be changed whether the switch is offline or online. If fmsmode is changed while the switch is online, a device reset is performed for the control device and an RSCN is generated with PID 0xDDFE00 (where 0xDD is the domain ID of the switch).

When FMS mode is on, the Fabric OS CLI commands listed here return a `switch busy` response if they are issued when the host-based management tool is performing a write operation. This serialization prevents interference from local switch commands when a host-based management program is being used to administer the switch.

<code>bladeDisable</code>	<code>slotOff</code>
<code>bladeEnable</code>	<code>slotOn</code>
<code>portDisable</code>	<code>switchCfgPersistentDisable</code>
<code>portEnable</code>	<code>switchDisable</code>
<code>portName</code>	<code>switchEnable</code>
<code>portShow</code>	<code>switchName</code>
<code>portSwap</code>	<code>switchShow</code>

 **NOTE:** You cannot use the `portCfgPersistentEnable` and `portCfgPersistentDisable` commands to persistently enable and disable ports when FMS mode is on. See the procedure [“Persistently enabling and disabling ports”](#) on page 115.

Changing `fmsmode` from `disabled` to `enabled` triggers the following events:

- Access to switch parameters is serialized.
- The active CUP configuration data is established as follows:
 - Port and switch names are not read from the IPL; they remain as previously set.
 - Port Block and Unblock values are not read from the IPL; they remain as previously set with the `portEnable` and `portDisable` commands.
 - Prohibit Dynamic Connectivity Mask (PDCM) values are read from the IPL; the default is Allow All.
- HP Advanced Zoning, if used, continues to be in force. If there are any differences in restrictions set up with HP Advanced Zoning and PDCM, the most restrictive rules are applied.
- RSCNs are sent to devices if PDCM results in changes to connectivity between a set of ports.

Changing `fmsmode` from `enabled` to `disabled` triggers the following events:

- A device reset is performed on the control device.
- PDCM is no longer enforced.
- RSCNs might be generated to some devices if PDCM removal results in changes to connectivity between a set of ports.
- If a given port was set to Block or Unblock, that port remains disabled or enabled.
- Serialized access to switch parameters ceases.

Setting up CUP when FMS mode is enabled

`Fmsmode` may be enabled and in use on a switch without a CUP license. The transition from `fmsmode disabled` to `fmsmode enabled` with the CUP license installed triggers the notification to the host systems that the CUP feature is available. Without this notification, the host systems never know the CUP feature is available, and consequently never try to communicate with it. Hence, it is possible that `fmsmode` may already be enabled on the switch.

If FMS mode is already enabled, set up CUP as follows:

1. Verify that FMS mode is enabled by entering the `ficoncupshow fmsmode` command.

If FMS mode is not enabled, see [“Enabling and disabling FMS mode”](#) on page 112.

△ **CAUTION:** If `fmsmode` is already enabled, disabling it might be disruptive to operation, because ports that were previously prevented from communicating are now able to do so.

2. If FMS mode is enabled, disable it by issuing the `ficoncupset fmsmode disable` command.
3. Install a CUP license key as described in [“Adding and removing FICON CUP licenses”](#) on page 116.
4. Issue the `ficoncupset fmsmode enable` command.

Displaying the `fmsmode` setting

The `ficoncupshow fmsmode` command displays the effective `fmsmode` setting for the switch. For example:

```
switch:admin> ficoncupshow fmsmode
fmsmode for the switch: Enabled
```

Displaying mode register bit settings

The mode register bits are described in [Table 19](#).

Table 19 FICON CUP mode register bits

Bit	Description
POSC	Programmed offline state control. When this bit is set on, the host is prevented from taking the switch offline. The default setting is 1 (on).
UAM	User alert mode. When this bit is set on, a warning is issued when an action is attempted that writes CUP parameters on the switch. The default setting is 0 (off).
ASM	Active=saved mode. When this bit is set on, all CUP configuration parameters are persistent, meaning that they are saved in nonvolatile storage in the IPL file that is applied upon a cold reboot or a power cycle. The default setting is 1 (on).
DCAM	Switch clock alert mode. When this bit is set on, a warning is issued when the <code>date</code> , <code>tsClockServer</code> , or <code>tsTimeZone</code> commands are entered to set the time and date on the switch. The default setting is 0 (off).
ACP	Alternate control prohibited. Because the Fabric OS CLI, Advanced Web Tools, and Fabric Manager are considered to be switch consoles, this bit has no effect on their operation. Attempts to set CUP parameters through SNMP are denied when this bit is set on. The default setting is 1 (on).
HCP	Host control prohibited. When this bit is set on, the host is not allowed to set CUP parameters. The default setting is 1 (on).

The `ficoncupshow modereg` command displays the mode register bit settings for the switch. A display of 0 indicates that the mode register bit is set off; 1 indicates that the bit is set on.

The command format is:

```
ficoncupshow modereg [bitname]
```

where *bitname* is one of the mode register bits described in [Table 19](#).

For example, to display all mode register bit settings for the switch:

```
switch:admin> ficoncupshow modereg
POSC  UAM  ASM  DCAM  ACP  HCP
-----
      1   0   1   0   1   1
```

For example, to display the mode register bit HCP for the switch:

```
switch:admin> ficoncupshow modereg HCP
HCP
1
```

Setting mode register bits

The `ficoncupset modereg` command sets the FICON CUP mode register bits for the local switch. Consider the following when changing mode register bits:

- As required by the CUP protocol, the UAM bit cannot be changed using this command.
- All mode register bits except UAM are saved across power on/off cycles; the UAM bit is reset to 0 following a power-on.
- Mode register bits can be changed when the switch is offline or online. If the ACP or HCP bits are changed when the switch is online, they may take effect any time between the completion of the current command and the end of the CCW command chain (or the next alternate manager operation).

The command format is:

```
ficoncupset modereg [bitname] 0 | 1
```

where:

<i>bitname</i>	Specifies one of the mode register bits described in Table 19 on page 114.
0	Specifies that the bit is off.
1	Specifies that the bit is on.

The following example sets the mode register bit HCP to off:

```
switch:admin> ficoncupset modereg HCP 0
Mode register bit HCP has been set to 0.
```

The following example sets the mode register bit ACP to on:

```
switch:admin> ficoncupset modereg ACP 1
Mode register bit ACP has been set to 1.
```

Persistently enabling and disabling ports

When `fmsmode` is enabled, you cannot use the `portCfgPersistentEnable` and `portCfgPersistentDisable` commands to persistently enable and disable ports. Instead, use this procedure:

1. Issue the following command to display the mode register bit settings:

```
ficoncupshow modereg
```

2. Verify that the ASM bit is set on (1).
3. If the ASM bit is set off (0), issue the following command to set it on:

```
ficoncupset modereg asm 1
```

4. Use the `portEnable` and `portDisable` commands to enable and disable ports as necessary. The ports remain enabled or disabled after a switch reboot.

In the following example, the ASM bit is set to on, and then the port at slot 1, port 1 is enabled persistently:

```
switch:admin> ficoncupshow modereg
POSC  UAM  ASM  DCAM  ACP  HCP
-----
    1    0    0    0    1    1


switch:admin> ficoncupset modereg ASM 1
Mode register bit ASM has been set to 1.

switch:admin> portenable 1/1
```

Port and switch naming standards

Fabric OS handles differences in port and switch naming rules between CUP and itself as follows:

- CUP employs 8-bit characters in port address names and switch names; Fabric OS employs 7-bit characters. When `fmsmode` is enabled, all characters greater than 0x40 and not equal to 0xFF (EBCDIC code page 37 [0x25]) are allowed in the name; therefore, it is possible for a channel to set a name with non-printable characters. If a name contains non-printable characters, they are displayed as dots. The following characters are also displayed as dots: semicolon (;), comma (,), equal sign (=), and at sign (@).

 **NOTE:** Configuration files that contain non-printable characters should not be edited manually, because many editors replace non-printable characters with some other characters without warning the user first.

- CUP has a 24-character unique port name limitation; Fabric OS supports port names up to 32 characters long. When `fmsmode` is enabled, names longer than 24 characters are truncated.
- To ensure that port names are unique, the characters ~00, ~01, ~02, and so on are appended to them.
- CUP allows a 24-character switch name; Fabric OS limits the switch name to 15 characters. To reconcile this difference, Fabric OS files the first 15 characters in the WWN record and stores the extra characters for CUP use.

Adding and removing FICON CUP licenses

If `fmsmode` is enabled when the FICON CUP license is removed, the control device is reset. PDCM enforcement continues. If `fmsmode` is disabled when the FICON CUP license is removed, no special action is taken.

If `fmsmode` is enabled on a switch that does not have a FICON CUP license and then the license is installed, you must first disable and then reenable `fmsmode`. If `fmsmode` is disabled and a FICON CUP license is installed, no special action is required.

Zoning and PDCM considerations

The FICON PDCM controls whether communication between a pair of ports in the switch is prohibited or allowed. If there are any differences in restrictions set up with HP Advanced Zoning and PDCM, the most restrictive rules are applied.

All FICON devices should be configured in a single zone using the Domain, Area notation. PDCM can then be used to allow or prohibit access between specific port pairs.

PDCM persists across a failover because it is replicated at all times to the standby CP blade. The active PDCM configuration is saved to the IPL if the ASM bit is set on.

Backing up and restoring configurations

The Fabric OS `configUpload` command saves up to 16 FICON configuration files, including IPL files. For details on the behavior of the `configDownload` command, see ["Restoring configurations in a FICON environment"](#) on page 75.

Troubleshooting

The following sources provide useful problem-solving information:

- The standard support commands (`portLogDump`, `supportSave`, `supportShow`) or the Fabric Manager Event Log

By default, the FICON group in the `supportShow` output is disabled. To enable the capture of FICON data in the `supportShow` output, issue the `supportshowcfgenable ficon` command. After you get confirmation that the configuration has been updated, the following information is collected and appears in the output for the `supportShow` command:

- `ficoncupshow fmsmode`
- `ficoncupshow modereg`

- `ficonDbg dump rnid`
- `ficonDbg log`
- `ficonShow ilir`
- `ficonShow lirr`
- `ficonShow rlir`
- `ficonShow rnid`
- `ficonShow switchrnid`
- `ficucmd dump -A`
- Other detailed information for protocol-specific problems:
 - Display port data structures using the `ptDataShow` command.
 - Display port registers using the `ptRegShow` command.

Identifying ports

The `ficonshow rlir` command displays, among other information, a tag field for the switch port. You can use this tag to identify the port on which a FICON link incident occurred. The tag field is a concatenation of the switch domain ID and port number, in hexadecimal format. The following example shows a link incident for the switch port at domain ID 120, port 93 (785D in hexadecimal):

```
switch:admin> ficonshow rlir
{
  {Fmt  Type PID      Port      Incident Count  TS Format      Time Stamp
   0x18 F    785d00   93              1 Time server Thu Apr 22 09:13:32 2004
   Port Status:          Link not operational
   Link Failure Type:    Loss of signal or synchronization

   Registered Port WWN      Registered Node WWN      Flag  Node Parameters
   50:05:07:64:01:40:16:03  50:05:07:64:00:c1:69:ca  0x10  0x200115
   Type number:             002064
   Model number:            103
   Manufacturer:            IBM
   Plant of Manufacture:    02
   Sequence Number:         0000000169CA
   tag:                     155d

   Switch Port WWN          Switch Node WWN          Flag  Node Parameters
   20:5d:00:60:69:80:45:7c  10:00:00:60:69:80:45:7c  0x00  0x200a5d
   Type number:             SLKWRM
   Model number:            24K
   Manufacturer:            BRD
   Plant of Manufacture:    CA
   Sequence Number:         000000000078
   tag:                     785d
}
}
The Local RLIR database has 1 entry.
switch:admin> ficonshow rlir
```

Backing up FICON files

The FICON file access facility is used to store configuration files. This includes IPL and other configuration files. Fabric OS saves the IPL and all other configuration files on the switch. A maximum of 16 configuration files, including the IPL file, are supported.

You can upload the configuration files saved on the switch to a management workstation using the `configUpload` command. If the switch loses the configuration due to a hardware failure or filesystem error, use the `configDownload` command to restore previously uploaded configuration files. Because data uploaded using the `configUpload` command also contains the IPL, if Active=Saved mode is enabled, the switch ignores the IPL file downloaded via the `configDownload` command.

Uploading the configuration files

Issue the `configUpload` command.

When you execute the `configUpload` command, all the files saved in the file access facility are uploaded to a management workstation. (There is a section in the uploaded configuration file labeled `FICON_CUP` that exists in an encoded format.)

Downloading configuration files with Active=Saved mode enabled

Issue the `configDownload` command.

The contents of existing files saved on the switch, which are also present in the `FICON_CUP` section, are overwritten.

The files in the `FICON` section of the configuration file, which are not currently on the switch, are saved on the switch.

The IPL is not replaced because Active=Saved mode is enabled. A warning message is displayed in the event log to warn users that the IPL is not overwritten.

Downloading configuration files with Active=Saved mode disabled

Issue the `configDownload` command.

The contents of existing files saved on the switch, which are also present in the `FICON_CUP` section, are overwritten.

The files in the `FICON` section of configuration file, which are not currently on the switch, are saved on the switch.

The IPL is replaced because Active=Saved mode is disabled.

Recording configuration information

You can use the following worksheet for recording FICON configuration information.

Sample IOCP configuration file for the SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128

The channel subsystem controls communication between a configured channel, the CU, and the device. The IOCDS defines the channels, CUs, and devices to the designated logical partitions (LPARs) within the server; this is defined using the Input/Output Configuration Program (IOCP). The IOCP statements are typically built using the hardware configuration dialog box (HCD). The interactive dialog box is used to generate your Input/Output Definition File (IODF), invoke the IOCP program, and subsequently build your production IOCDS.

Each FICON director in a fabric must have a unique domain ID and a unique switch ID. The switch ID used in the IOCP definitions can be any value between x00 to xFF. The domain ID range for directors is hex x01 to xEF or decimal 1 to 239. When defining the switch IDs in the IOCP definitions, ensure that you use values within the FICON director's range.

The switch ID has to be assigned by the user and must be unique within the scope of the definitions (IOCP and HCD).

The domain ID is assigned by the manufacturer and can be customized to a different value. It must be unique within the fabric.

HP recommends that the switch ID (in IOCP or HCD) be set to the same value as the domain ID of the FICON director, which is defined to the FICON director at installation time. This simplifies the configuration and reduces confusion by having a common designation across all definitions.

For more information, see the IBM publication *zSeries Input/Output Configuration Program User's Guide for ICP IOCP* (SB10-7037).

In the following sample IOCP configuration file, the UNIT value for FICON CUP definitions is 2032 for any FICON director regardless of vendor or platform. All Domain IDs are specified in hex values in the

IOCP (and not in decimal values); the Domain IDs in the following example are for demonstration purposes only.

```
*-----
* SilkWorm 24000 Domain_ID=61 (in hex)
*-----
CNTLUNIT CUNUMBR=0D8,UNITADD=00,UNIT=2032,
        PATH=( 50,51),
        LINK=( 61FE,61FE)
IODEVICE ADDRESS=( 0D8,1),CUNUMBR=0D8,UNIT=2032,STADET=Y,UNITADD=00
*-----
*SilkWorm 12000 Domain_ID=22 (left side logical switch 0 in hex)
*-----
CNTLUNIT CUNUMBR=0D9,UNITADD=00,UNIT=2032,
        PATH=( 8A,8B),
        LINK=( 22FE,22FE)
IODEVICE ADDRESS=( 0D9,1),CUNUMBR=0D9,UNIT=2032,STADET=Y,UNITADD=00
*-----
* SilkWorm 12000 Domain_ID=23 (right side logical switch 1 in hex)
*-----
*
CNTLUNIT CUNUMBR=0DA,UNITADD=00,UNIT=2032,
        PATH=( 92,93),
        LINK=( 23FE,23FE)
IODEVICE ADDRESS=( 0DA,1),CUNUMBR=0DA,UNIT=2032,STADET=Y,UNITADD=00
*-----
* SilkWorm 3900 Domain_ID=25 (in hex)
*-----
*
CNTLUNIT CUNUMBR=0DB,UNITADD=00,UNIT=2032,
        PATH=( 5A,5B),
        LINK=( 25FE,25FE)
IODEVICE ADDRESS=( 0DB,1),CUNUMBR=0DB,UNIT=2032,STADET=Y,UNITADD=00
*
*-----
```

Sample Resource Management Facility configuration file for mainframe

Resource management facility (RMF) is a performance management tool that measures selected areas of system activity, including different views of the FICON channel. RMF presents data collected in the form of System Management Facility (SMF) records. This data is essential for any kind of FICON channel performance troubleshooting.

To obtain an RMF FICON director activity report, you must include the keyword **FCD** in the RMF configuration file for the FICON director (this is generic for any FICON director). You must also define the CUP port. In the following example, the keyword is boldfaced.

```

/*****
/*  MONITOR I OPTIONS
/*
/*          X A      O N L Y
/*
/*
*****/

      FCD                      /* FICON Director */
CHAN                      /* COLLECT CHANNEL STATISTICS */
CPU                      /* COLLECT CPU STATISTICS */
CYCLE(1000)              /* SAMPLE ONCE EVERY SECOND */
DEVICE(NOSG)             /* PREVENT SORT OF STORAGE GROUPS*/
DEVICE(NOCHRDR)          /* CHARACTER READER STATISTICS
                          WILL NOT BE COLLECTED */
DEVICE(COMM)             /* COMMUNICATION EQUIPMENT STATS.
                          WILL BE COLLECTED PDS 1/25/94 */
DEVICE(DASD)             /* DIRECT ACCESS DEVICE STATISTICS
                          WILL BE COLLECTED */
DEVICE(GRAPH)            /* GRAPHICS DEVICE STATISTICS
                          WILL BE COLLECTED */
DEVICE(TAPE)             /* TAPE DEVICE STATISTICS
                          WILL BE COLLECTED */
DEVICE(NOUNITR)          /* UNIT RECORD DEVICE STATISTICS
                          WILL NOT BE COLLECTED */
DEVICE(NONMBR)           /* NO DEVICE SELECTIVITY BY
                          DEVICE NUMBERS */
IOQ(DASD)               /* COLLECT DASD I/O QUEUING
                          STATISTICS */
IOQ(NOCHRDR)            /* PREVENT CHARACTER READER
                          I/O QUEUING STATISTICS */
IOQ(NOUNITR)            /* PREVENT UNIT RECORD DEVICE
                          I/O QUEUING STATISTICS */
IOQ(COMM)               /* COLLECT COMMUNICATION EQUIPMENT
                          I/O QUEUING STATS PDS 1/25/94 */
IOQ(GRAPH)              /* COLLECT GRAPHICS DEVICE

```

8 Configuring the Distributed Management Server

The HP Fabric OS Distributed Management Server allows a SAN management application to retrieve information and administer interconnected switches, servers, and storage devices. The management server assists in the autodiscovery of switch-based fabrics and their associated topologies.

A client of the Distributed Management Server can find basic information about the switches in the fabric and use this information to construct topology relationships. The management server also allows you to obtain certain switch attributes and, in some cases, modify them. For example, logical names identifying switches can be registered with the management server.

The management server provides several advantages for managing a Fibre Channel fabric:

- It is accessed by an external Fibre Channel node at the well-known address FFFFFAh, so an application can access information about fabric management with minimal knowledge of the existing configuration.
- It is replicated on every HP StorageWorks switch within a fabric.
- It provides an unzoned view of the overall fabric configuration. This fabric topology view exposes the internal configuration of a fabric for management purposes; it contains interconnect information about switches and devices connected to the fabric. Under normal circumstances, a device (typically an FCP initiator) queries the Name Server for storage devices within its member zones. Because this limited view is not always sufficient, the management server provides the application with a list of the entire Name Server database.

 **NOTE:** The management server platform service is available only with Fabric OS 2.3.0 and later.

Enabling and disabling the platform services

The management server is located at the Fibre Channel well-known address FFFFFAh. All management server services except platform services are enabled by default.

Enabling platform services

1. Connect to the switch and log in as admin.
2. Issue the `msplMgmtActivate` command.

For example:

```
switch:admin> msplmgmtactivate
Request to activate MS Platform Service in progress.....
*Completed activating MS Platform Service in the fabric!
switch:admin>
```

Disabling platform services

1. Connect to the switch and log in as admin.
2. Issue the `msplMgmtDeactivate` command.
3. Enter `y` to confirm the deactivation.

For example:

```
switch:admin> msplmgmtdeactivate
MS Platform Service is currently enabled.
This will erase MS Platform Service configuration
information as well as database in the entire fabric.
Would you like to continue this operation? (yes, y, no, n): [no] y
Request to deactivate MS Platform Service in progress.....
*Completed deactivating MS Platform Service in the fabric!
switch:admin>
```


Controlling access

You can use the `msConfigure` command to control access to the management server database.

An ACL of WWN addresses determines which systems have access to the management server database. The ACL typically contains those WWNs of host systems that are running management applications.

If the list is empty (the default), the management server is accessible to all systems connected in-band to the fabric. For more access security, you can specify WWNs in the ACL so that access to the management server is restricted only to those WWNs listed.

The ACL is switch-based. Therefore, only hosts that are connected directly to the switch are affected by the ACL. A host that is somewhere else in the fabric and is connected to a switch with an empty ACL is allowed to access the management server.

 **NOTE:** The `msConfigure` command is disabled if the switch is in secure mode. See the *HP StorageWorks Secure Fabric OS administrator guide* for more information.

Displaying the management server ACL

1. Connect to the switch and log in as admin.
2. Issue the `msConfigure` command.
The command becomes interactive.
3. At the select prompt, enter 1 to display the access list.
A list of WWNs that have access to the management server is displayed.

In the following example, the list is empty:

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 1
MS Access list is empty.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
done ...
switch:admin>
```

Adding a member to the ACL

1. Connect to the switch and log in as admin.
2. Issue the `msConfigure` command.
The command becomes interactive.

3. At the select prompt, enter 2 to add a member based on its port/node WWN.
4. Enter the WWN of the host to be added to the ACL.
5. At the select prompt, enter 1 to verify the WWN you entered was added to the ACL.
6. After verifying that the WWN was added correctly, enter 0 at the prompt to end the session.
7. At the Update the FLASH? prompt, enter y.
8. Press **Enter** to update the nonvolatile memory and end the session.

For example:

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 2
Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully added to the MS ACL.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (14): {
  20:00:00:20:37:65:ce:aa
  20:00:00:20:37:65:ce:bb
  20:00:00:20:37:65:ce:ff
  20:00:00:20:37:65:ce:11
  20:00:00:20:37:65:ce:22
  20:00:00:20:37:65:ce:33
  20:00:00:20:37:65:ce:44
  10:00:00:60:69:04:11:24
  10:00:00:60:69:04:11:23
  21:00:00:e0:8b:04:70:3b
  10:00:00:60:69:04:11:33
  20:00:00:20:37:65:ce:55
  20:00:00:20:37:65:ce:66
  00:00:00:00:00:00:00:00
}
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH? (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.
switch:admin>
```

Deleting a member from the ACL

1. Connect to the switch and log in as admin.
2. Issue the msConfigure command.
The command becomes interactive.
3. At the select prompt, enter 3 to delete a member based on its port/node WWN.
4. At the select prompt, enter the WWN of the member to be deleted from the ACL.
5. At the select prompt, enter 1 to verify the WWN you entered was deleted from the ACL.
6. After verifying that the WWN was deleted correctly, enter 0 at the prompt to end the session.

7. At the Update the FLASH? prompt, enter y.
8. Press **Enter** to update the nonvolatile memory and end the session.

For example:

```
switch:admin> msconfigure
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 3
Port/Node WWN (in hex): [00:00:00:00:00:00:00:00] 20:00:00:20:37:65:ce:aa
*WWN is successfully deleted from the MS ACL.
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [2] 1
MS Access List consists of (13): {
  20:00:00:20:37:65:ce:aa
  20:00:00:20:37:65:ce:bb
  20:00:00:20:37:65:ce:ff
  20:00:00:20:37:65:ce:11
  20:00:00:20:37:65:ce:22
  20:00:00:20:37:65:ce:33
  10:00:00:60:69:04:11:24
  10:00:00:60:69:04:11:23
  21:00:00:e0:8b:04:70:3b
  10:00:00:60:69:04:11:33
  20:00:00:20:37:65:ce:55
  20:00:00:20:37:65:ce:66
}
0      Done
1      Display the access list
2      Add member based on its Port/Node WWN
3      Delete member based on its Port/Node WWN
select : (0..3) [1] 0
done ...
Update the FLASH? (yes, y, no, n): [yes] y
*Successfully saved the MS ACL to the flash.
switch:admin>
```

Configuring the server database

The management server database can be viewed or cleared.

Viewing the contents of the management server database

1. Connect to the switch and log in as admin.
2. Issue the msPlatShow command.

The contents of the management server platform database are displayed. For example:

```
switch:admin> msplatshow
-----
Platform Name: [9] "first obj"
Platform Type: 5 : GATEWAY
Number of Associated M.A.: 1
[35] "http://java.sun.com/products/plugin"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:71
-----
Platform Name: [10] "second obj"
Platform Type: 7 : HOST_BUS_ADAPTER
Number of Associated M.A.: 1
Associated Management Addresses:
[30] "http://java.sun.com/products/1"
Number of Associated Node Names: 1
Associated Node Names:
10:00:00:60:69:20:15:75
```

Clearing the management server database

1. Connect to the switch and log in as admin.
2. Issue the `msplClearDb` command.
3. Enter `y` to confirm the deletion.

The management server platform database is cleared.

Controlling topology discovery

The topology discovery feature can be displayed, enabled, and disabled; it is disabled by default.

Displaying topology discovery status

1. Connect to the switch and log in as admin.
2. Issue the `mstdReadConfig` command.

For example:

```
switch:admin> mstdreadconfig
*MS Topology Discovery is Enabled.
switch:admin>
```

Enabling topology discovery

1. Connect to the switch and log in as admin.
2. Issue the `mstdEnable` command to enable the discovery feature locally or issue the `mstdEnable all` command to enable the discovery feature on the entire fabric.

For example:

```
switch:admin> mstdenable

Request to enable MS Topology Discovery Service in progress....
*MS Topology Discovery enabled locally.

switch:admin> mstdenable ALL

Request to enable MS Topology Discovery Service in progress....
*MS Topology Discovery enabled locally.
*MS Topology Discovery Enable Operation Complete!!
```

Disabling topology discovery

1. Connect to the switch and log in as admin.
2. Issue the `mstdDisable` command to disable the discovery feature locally.
A warning is displayed, informing you that all NID entries might be cleared.
3. Enter `y` to disable the discovery feature.
4. Issue the `mstdDisable all` command to disable the discovery feature on the entire fabric.
5. Enter `y` to complete the disabling of the discovery feature.

For example:

```
switch:admin> mstddisable
This may erase all NID entries. Are you sure? (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.

switch:admin> mstddisable all
This may erase all NID entries. Are you sure? (yes, y, no, n): [no] y

Request to disable MS Topology Discovery Service in progress....
*MS Topology Discovery disabled locally.
*MS Topology Discovery Disable Operation Complete!!
```

9 Working with diagnostic features

This chapter provides information on diagnostics and how to display system, port, and specific hardware information. It also describes how to set up system logging mapping (`syslogd`) and how to set up the offloading of error messages (`supportSave`).

The purpose of the diagnostic subsystem is to evaluate the integrity of the system hardware.

Diagnostics are invoked two ways:

- Automatically during the POST
- Manually using Fabric OS CLI commands

The error messages generated during these test activities are sent to the serial console and system message logs, whose output formats may differ slightly.

Use the `diagHelp` command to receive a list of all available diagnostic commands.

See the *HP StorageWorks Fabric OS 5.x command reference guide* for a complete description of each command.

Viewing POST

By default, when you power on the system, the boot loader performs POST and loads a Fabric OS kernel image.

POST provides a quick indication of hardware readiness when hardware is powered up. POST does not require user input to function. It usually completes within several minutes, and supports minimal validation because of the restriction on test duration. POST performs a basic health check before a new switch joins a fabric.

POST consists of two groups: POST1 and POST2. POST1 validates the hardware interconnect of the device, and POST2 validates the ability of the device to pass data frames between the ports. The specific set of diagnostic and test commands run during POST depends on the switch model.

POST1 cannot be bypassed and runs from the boot loader. The factory default configuration is also set to run POST2, but you can configure your switch to bypass POST2, which runs after the kernel image has started but before general system services, such as login, are enabled.

Although each test performed during POST2 is configurable, modify a POST2 test only if directed by your switch provider's customer service representative.

You can use the `diagDisablePost` command to disable POST2, and you can reenable it using the `diagEnablePost` command. See the *HP StorageWorks Fabric OS 5.x command reference guide* for additional information about these commands.

The following example shows a typical boot sequence, including POST messages:

```
The system is coming up, please wait...

Read board ID of 0x80 from addr 0x23
Read extended model ID of 0x16 from addr 0x22
Matched board/model ID to platform index 4
PCI Bus scan at bus 0
:   :   :
:   :   :
Checking system RAM - press any key to stop test

Checking memory address: 00100000

System RAM test using Default POST RAM Test succeeded.

Press escape within 4 seconds to enter boot interface.
Booting "Fabric Operating System" image.

Linux/PPC load:
BootROM command line: quiet
Uncompressing Linux...done.
Now booting the kernel
Attempting to find a root file system on hda2...
modprobe: modprobe: Can't open dependencies file /lib/modules/2.4.19/modules.dep (No such file or directory)
INIT: version 2.78 booting
INIT: Entering runlevel: 3
eth0: Link status change: Link Up. 100 Mbps Full duplex Auto (autonegotiation complete).

INITCP: CPLD Vers: 0x95 Image ID: 0x19
uptime: 2008; sysc qid: 0
Fabric OS (Paulsa45)
Paulsa45 console login: 2005/03/31-20:12:42, [TRCE-5000], 0,, INFO, ?, trace:, trace_buffer.c, line: 1170

2005/03/31-20:12:42, [LOG-5000], 0,, INFO, SW4100_P45, Previous message repeat 1 time(s), trace_ulib.c, line: 540
2005/03/31-20:12:43, [HAM-1004], 219,, INFO, SW4100_P45, Processor rebooted - Unknown
SNMP Research SNMP Agent Resident Module Version 15.3.1.4
Copyright 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001 SNMP Research, Inc.
sysctrl: all services Standby
FSSK 2: chassis0(0): state not synchronized
FSSK 2: Services starting a COLD recovery
2005/03/31-20:12:48, [FSS-5002], 0,, INFO, SW4100_P45, chassis0(0): state not synchronized, svc.c, line: 318
2005/03/31-20:12:48, [FSS-5002], 0,, INFO, SW4100_P45, Services starting a COLD recovery, mdev.c, line: 638
2005/03/31-20:12:49, [MFIC-1002], 220,, INFO, Paulsa45, Chassis FRU header not programmed for switch NID, using
defaults (applies only to FICON environments).
sysctrl: all services Active
2005/03/31-20:12:50, [DGD-5001], 0,, INFO, SW4100_P45, Slot 0 has started POST., main.c, line: 1189
POST1: Started running Thu Mar 31 20:12:51 GMT 2005
POST1: Test #1 - Running turboramtest
POST1: Test #2 - Running portregtest
POST1: Script PASSED with exit status of 0 Thu Mar 31 20:12:54 GMT 2005 took (0:0:3)
POST2: Started running Thu Mar 31 20:12:55 GMT 2005
POST2: Test #1 - Running portloopbacktest (SERDES)
POST2: Test #2 - Running minicycle (SERDES)
POST2: Running diagshow
POST2: Script PASSED with exit status of 0 Thu Mar 31 20:13:12 GMT 2005 took (0:0:17)
2005/03/31-20:13:13, [BL-1000], 221,, INFO, Paulsa45, Initializing Ports... Enabling switch...
2005/03/31-20:13:13, [BL-1001], 222,, INFO, Paulsa45, Port Initialization Completed
2005/03/31-20:13:13, [EM-5012], 0,, INFO, SW4100_P45, EM: sent dumpready to ME., em.c, line: 2152
2005/03/31-20:13:13, [DGD-5002], 0,, INFO, SW4100_P45, Slot 0 has passed the POST tests., main.c, line: 936
```

If you choose to bypass POST2, or after POST2 completes, various system services are started and the boot process displays additional console status and progress messages.

Viewing switch status

Use the `switchStatusShow` command to display the overall status of the switch, including its power supplies, fans, and temperature. If the status of any one of these components is either marginal or down, the overall status of the switch is also displayed as marginal or down. If all components have a healthy status, the switch displays a healthy status.

To modify the rules used to classify the health of each component, use the `switchStatusPolicySet` command. To view the rules, use the `switchStatusPolicyShow` command.

Viewing the overall status of the switch

1. Connect to the switch and log in as admin.
2. Issue the `switchStatusShow` command.

For example:

```
switch:admin> switchstatusshow
Switch Health Report Report time: 03/21/2005 03:50:36 PM
Switch Name:      SW3900
IP address:       10.33.54.176
SwitchState:      MARGINAL
Duration:         863:23
Power supplies monitor MARGINAL
Temperatures monitor HEALTHY
Fans monitor      HEALTHY
WWN monitor       HEALTHY
Standby CP monitor HEALTHY
Blades monitor    HEALTHY
Flash monitor     HEALTHY
Marginal ports monitor HEALTHY
Faulty ports monitor HEALTHY
Missing SFPs monitor HEALTHY
All ports are healthy
switch:admin>
```

For more information on how the overall switch status is determined, see the `switchStatusPolicySet` command in the *HP StorageWorks Fabric OS 5.x command reference guide*.

Displaying switch information

1. Connect to the switch and log in as admin.
2. Issue the `switchShow` command. This command displays the following information for a switch:

- `switchname` displays the switch name.
- `switchtype` displays the switch model and firmware version numbers.
- `switchstate` displays the switch state: Online, Offline, Testing, or Faulty.
- `switchrole` displays the switch role: Principal, Subordinate, or Disabled.
- `switchdomain` displays the switch Domain ID.
- `switchid` displays the embedded port D_ID of the switch.
- `switchwwn` displays the switch WWN.
- `switchbeacon` displays the switch beaconing state: either ON or OFF.

The `switchShow` command also displays the following information for ports on the specified switch:

- Module type: The SFP type, if an SFP is present.
- Port speed: The speed of the Port (1G, 2G, 4G, N1, N2, N4, or AN). The speed can be fixed, negotiated, or auto-negotiated.
- Port state: The port status.
- Comment: Information about the port. This section might be blank or display the WWN for the F_Port or E_Port, trunking state, upstream or downstream status.

The details displayed for each switch differ for different switch models. For more information, see the `switchShow` command in the *HP StorageWorks Fabric OS 5.x command reference guide*.

Displaying the uptime for a switch

1. Connect to the switch and log in as admin.
2. Issue the `uptime` command.

For example:

```
switch:admin> uptime
4:43am up 1 day, 12:32, 1 user, load average: 1.29, 1.31, 1.27
switch:admin>
```

The `uptime` command displays the length of time the system has been in operation, the total cumulative amount of uptime since the system was first powered-on, the date and time of the last reboot (applies only to Fabric OS 3.x and 2.6.x systems), the reason for the last reboot (applies only to Fabric OS 3.x and 2.6.x systems), and the load average over the past one minute (1.29 in the preceding example), five minutes (1.31 in the example), and 15 minutes (1.27 in the example). The reason for the last switch reboot is also recorded in the system message log.

Viewing port information

Use the commands that follow to view information about ports.

Viewing the status of a port

1. Connect to the switch and log in as admin.
2. Issue the `portShow` command, specifying the number that corresponds to the port you are troubleshooting.

In the following example, the status of port two is shown:

```
switch:admin> portshow 2
portName:
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x4903      PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE LOGIN LED
portType: 10.0
portState: 1 Online
portPhys: 6 In_Sync
portScn: 16 E_Port Trunk port
port generation number: 351
portId: 290200
portIfId: 43020004
portWwn: 20:02:00:05:1e:34:01:be
portWwn of device(s) connected:
None
Distance: normal
portSpeed: N4Gbps

Interrupts:      0      Link_failure: 50      Frjt:      0
Unknown:         0      Loss_of_sync: 55      Fbsy:      0
Lli:             524     Loss_of_sig: 54
Proc_rqrd:       0      Protocol_err: 0
Timed_out:       0      Invalid_word: 0
Rx_flushed:      0      Invalid_crc: 0
Tx_unavail:      0      Delim_err: 0
Free_buffer:     0      Address_err: 0
Overrun:         0      Lr_in:      100
Suspended:       0      Lr_out:     50
Parity_err:      0      Ols_in:     50
2_parity_err:    0      Ols_out:    52
CMI_bus_err:     0

switch:admin>
```

See the *HP StorageWorks Fabric OS 5.x command reference guide* for additional `portShow` command information, such as the syntax for slot or port numbering.

Displaying the port statistics

1. Connect to the switch and log in as admin.
2. Issue the `portStatsShow` command.

Port statistics include information such as number of frames received, number of frames sent, number of encoding errors received, and number of class 2 and class 3 frames received.

See the *HP StorageWorks Fabric OS 5.x command reference guide* for additional `portStatsShow` command information, such as the syntax for slot or port numbering. The following example shows the output from the `portStatsShow` command:

```
switch:admin> portstatsshow 3/7
stat_wtx          0          4-byte words transmitted
stat_wrx          0          4-byte words received
stat_ftx          0          Frames transmitted
stat_frx          0          Frames received
stat_c2_frx       0          Class 2 frames received
stat_c3_frx       0          Class 3 frames received
stat_lc_rx        0          Link control frames received
stat_mc_rx        0          Multicast frames received
stat_mc_to        0          Multicast timeouts
stat_mc_tx        0          Multicast frames transmitted
tim_rdy_pri       0          Time R_RDY high priority
tim_txcrd_z       0          Time BB credit zero
er_enc_in         0          Encoding errors inside of frames
er_crc            0          Frames with CRC errors
er_trunc          0          Frames shorter than minimum
er_toolong        0          Frames longer than maximum
er_bad_eof        0          Frames with bad end-of-frame
er_enc_out        0          Encoding error outside of frames
er_bad_os         0          Invalid ordered set
er_c3_timeout     0          Class 3 frames discarded due to timeout
er_c3_dest_unreach 0          Class 3 frames discarded due to destination unreachable
er_other_discard  0          Other discards
er_crc_good_eof   0          Crc error with good eof
er_inv_arb        0          Invalid ARB
open              0          loop_open
transfer          0          loop_transfer
opened           0          FL_Port opened
starve_stop       0          tenancies stopped due to starvation
fl_tenancy        0          number of times FL has the tenancy
nl_tenancy        0          number of times NL has the tenancy
zero_tenancy      0          zero tenancy

switch:admin>
```

Displaying a summary of port errors for a switch

1. Connect to the switch and log in as admin.
2. Issue the `portErrShow` command.

See the *HP StorageWorks Fabric OS 5.x command reference guide* for additional `portErrShow` command information.

The following example shows output from the `portErrShow` command:

```
switch:admin> porterrshow
      frames  enc  crc  too  too  bad  enc  disc  link  loss  loss  frjt  fbsy
      tx   rx   in  err shrt long  eof  out   c3 fail sync sig
sig=====
0:   22   24   0   0   0   0   0  1.5m   0   7   3   0   0   0
1:   22   24   0   0   0   0   0  1.2m   0   7   3   0   0   0
2:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
3:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
4:  149m  99m   0   0   0   0   0  448    0   7   6   0   0   0
5:  149m  99m   0   0   0   0   0  395    0   7   6   0   0   0
6:  147m  99m   0   0   0   0   0  706    0   7   6   0   0   0
7:  150m  99m   0   0   0   0   0  160    0   7   5   0   0   0
8:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
9:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
10:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
11:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
12:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
13:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
14:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
15:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
32:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
33:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
34:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
35:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
36:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
37:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
38:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
39:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
40:  99m 146m   0   0   0   0   0  666    0   6  796   7   0   0
41:  99m 149m   0   0   0   0   0  15k    0   2  303   4   0   0
42:  99m 152m   0   0   0   0   0  665    0   2  221   5   0   0
43:  99m 147m   0   0   0   0   0  16k    0   2  144   4   0   0
44:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
45:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
46:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
47:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
```

The `portErrShow` command output provides one output line per port. See [Table 20](#) for a description of the error types.

Table 20 Port error summary description


Error type	Description
frames tx	Frames transmitted
frames rx	Frames received
enc in	Encoding errors inside frames
crc err	Frames with cyclic redundancy check (CRC) errors
too shrt	Frames shorter than minimum
too long	Frames longer than maximum
bad eof	Frames with bad end-of-frame delimiters
enc out	Encoding error outside of frames
disc c3	Class 3 frames discarded
link fail	Link failures (LF1 or LF2 states)
loss sync	Loss of synchronization
loss sig	Loss of signal

Table 20 Port error summary description (continued)

Error type	Description
frjt	Frames rejected with F_RJT
fbsy	Frames busied with F_BSY

Viewing equipment status

You can display status for fans, power supply, and temperature.

 **NOTE:** The number of fans, power supply units, and temperature sensors depends on the switch type. For detailed specifications on these components, see the SAN Switch installation guide for your switch model. The specific output from the status commands varies, depending on the switch type.

Displaying the status of the fans

1. Connect to the switch and log in as admin.
2. Issue the `fanShow` command.

For example:

```
switch:admin> fanshow
Fan 1 is OK           speed is 7010 RPM
Fan 2 is OK           speed is 7180 RPM
Fan 3 is OK           speed is 7068 RPM
Fan 4 is OK           speed is 7116 RPM
Fan 5 is OK           speed is 7155 RPM
Fan 6 is OK           speed is 7001 RPM
switch:admin>
```

The possible status values are:

- OK: Fan is functioning correctly.
- Absent: Fan is not present.
- Below minimum: Fan is present but rotating too slowly or stopped.
- Above minimum: Fan is rotating too quickly.
- Unknown: An unknown fan unit is installed.
- FAULTY: Fan has exceeded hardware tolerance.

Displaying the status of a power supply

1. Connect to the switch and log in as admin.
2. Issue the `psShow` command.

For example:

```
switch:admin> psshow
Power Supply #1 is OK
0335,FF2Z0007161,60-0000739-02, B,,DCJ3002-01P, B,FF2Z0007161
Power Supply #2 is faulty
0335,FF2Z0007176,60-0000739-02, B,,DCJ3002-01P, B,FF2Z0007176
switch:admin>
```

The possible status values are:

- OK: Power supply is functioning correctly.
- Absent: Power supply is not present.
- Unknown: An unknown power supply unit is installed.
- Predicting failure: Power supply is present but predicting failure.
- FAULTY: Power supply is present but faulty (no power cable, power switch turned off, fuse blown, or other internal error).

Displaying temperature status

1. Connect to the switch and log in as admin.
2. Issue the `tempShow` command.

For example:

```
switch:admin> tempshow
```

Index	Status	Centigrade	Fahrenheit
1	OK	21	70
2	OK	22	72
3	OK	29	84
4	OK	24	75
5	OK	25	77

```
switch:admin>
```

Information is displayed for each temperature sensor in the switch.

The possible temperature status values are:

- OK: Temperature is within acceptable range.
- FAIL: Temperature is outside of acceptable range. Damage might occur.

Viewing the system message log

The system message log feature enables messages to be saved across power cycles and reboots.

The Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director maintain an independent system message log for each of the two CP blades. For these models, configure `syslogd` to support chronological system message logs. For details, see ["Configuring for syslogd"](#) on page 138.

For details on error messages, see the *HP StorageWorks Fabric OS 5.x diagnostics and system error messages reference guide*.

Displaying the system message log, with no page breaks

1. Connect to the switch and log in as admin.
2. Issue the `errDump` command.

Displaying the system message log, one page at a time

1. Connect to the switch and log in as admin.
2. Issue the `errShow` command.


Clearing the system message log

1. Connect to the switch and log in as admin.
2. Issue the `errClear` command.

All switch and chassis events are removed.

Viewing the port log

Fabric OS maintains an internal log of all port activity. The port log stores entries for each port as a circular buffer. Each port has space to store 8000 log entries. When the log is full, the newest log entries overwrite the oldest log entries. Port logs are not persistent and are lost over power-cycles and reboots. If the port log is disabled, an error message is displayed.

 **NOTE:** Port log functionality is completely separate from the system message log. Port logs are typically used to troubleshoot device connections.

1. Connect to the switch and log in as admin.
2. Issue the `portLogShow` command:

```
switch:admin> portlogshow 12

time          task      event  port cmd  args
-----
Thu Apr 14 12:07:09 2005
12:07:09.350  PORT      Rx      0   40  02ffffffd,00ffffffd,0608ffff,14000000
12:07:09.350  PORT      Tx      0   0   c0ffffffd,00ffffffd,060807fc
12:07:10.812  PORT      Tx      0   40  02ffffffd,00ffffffd,07feffff,14000000
12:07:10.813  PORT      Rx      0   0   c0ffffffd,00ffffffd,07fe0627
12:07:19.492  PORT      Tx      4   40  02ffffffd,00ffffffd,0800ffff,14000000
12:07:19.492  PORT      Tx     22  40  02ffffffd,00ffffffd,0802ffff,14000000
12:07:19.493  PORT      Rx      4   0   c0ffffffd,00ffffffd,08009287
12:07:19.493  PORT      Tx     24  40  02ffffffd,00ffffffd,0804ffff,14000000
12:07:19.494  PORT      Tx     31  40  02ffffffd,00ffffffd,0806ffff,14000000
12:07:19.494  PORT      Rx     22   0   c0ffffffd,00ffffffd,0802928d
12:07:19.494  PORT      Rx     24   0   c0ffffffd,00ffffffd,080492a3
12:07:19.495  PORT      Rx     31   0   c0ffffffd,00ffffffd,080692a7
```

Use the commands summarized in [Table 21](#) to view and manage port logs. See the *HP StorageWorks Fabric OS 5.x diagnostics and system error messages reference guide* for additional information about these commands.

Table 21 Commands for port log management

Command	Description
<code>portLogClear</code>	Clear port logs for all or particular ports.
<code>portLogDisable</code>	Disable port logs for all or particular ports.
<code>portLogDump</code>	Display port logs for all or particular ports, without page breaks.
<code>portLogEnable</code>	Enable port logs for all or particular ports.
<code>portLogShow</code>	Display port logs for all or particular ports, with page breaks.

The `portLogDump` command output (trace) is a powerful tool that is used to troubleshoot fabric issues. The `portLogDump` output provides detailed information about the actions and communications within a fabric. By understanding the processes that are taking place in the fabric, issues can be identified and located.

The `portLogDump` command displays the port log, showing a portion of the Fibre Channel payload and header (FC-PH). The header contains control and addressing information associated with the frame. The payload contains the information being transported by the frame and is determined by the higher-level service or FC_4 upper-level protocol. There are many different payload formats based on the protocol.

Because a portLogDump output is long, a truncated example is presented:


```
switch:admin> portlogdump
task event port cmd args
-----
16:30:41.780 PORT Rx 9 40 02ffffffd,00ffffffd,0061ffff,14000000
16:30:41.780 PORT Tx 9 0 c0ffffffd,00ffffffd,0061030f
16:30:42.503 PORT Tx 9 40 02ffffffd,00ffffffd,0310ffff,14000000
16:30:42.505 PORT Rx 9 0 c0ffffffd,00ffffffd,03100062
16:31:00.464 PORT Rx 9 20 02fffc01,00fffc0,0063ffff,01000000
16:31:00.464 PORT Tx 9 0 c0fffc0,00fffc01,00630311
16:31:00.465 nsd ctin 9 fc 000104a0,0000007f
16:31:00.465 nsd ctout 9 fc 00038002,00000003,01fffc01
16:31:00.466 PORT Tx 9 356 03fffc0,00fffc01,00630311,01000000
16:31:00.474 PORT Rx 9 0 c0fffc01,00fffc0,00630311
16:31:01.844 PORT Tx 9 40 02ffffffd,00ffffffd,0312ffff,14000000
16:31:01.854 PORT Rx 9 0 c0ffffffd,00ffffffd,03120064
16:31:01.963 PORT Rx 9 40 02ffffffd,00ffffffd,0065ffff,14000000
16:31:01.963 PORT Tx 9 0 c0ffffffd,00ffffffd,00650313
16:31:14.726 INTR pstate 0 LF2
16:31:14.729 PORT scn 0 137 00000000,00000000,00000008
16:31:14.729 PORT scn 0 129 00000000,00000000,00000400
16:31:14.729 PORT scn 0 2 00010004,00000000,00000002
16:31:14.730 SPEE sn 0 ws 00000002,00000000,00000000
<output truncated>
```

Configuring for syslogd

The system logging daemon (`syslogd`) is a process on UNIX, Linux, and some Windows systems that reads and logs messages as specified by the system administrator.

Fabric OS can be configured to use a UNIX-style `syslogd` process to forward system events and error messages to log files on a remote host system.

The host system can be running UNIX, Linux, or any other operating system that supports the standard `syslogd` functionality.

 **NOTE:** Fabric OS releases earlier than 4.4.0 do not support UNIX local7 facilities; they support kern facilities. Starting with Fabric OS 4.4.0, kern facilities are no longer supported; UNIX local7 facilities are supported. The default facility level is 7.

Configuring for `syslogd` involves configuring the host, enabling `syslogd` on the HP StorageWorks switch, and, as an option, setting the facility level.

Configuring the host

Fabric OS supports a subset of UNIX-style message severities that default to the UNIX local7 facility. To configure the host, edit the `/etc/syslog.conf` file to map Fabric OS message severities to UNIX severities, as shown in [Table 22](#).

Table 22 Fabric OS and UNIX message severities

Fabric OS message severity	UNIX message severity
Critical (1)	Emergency (0)
Error (2)	Error (3)
Warning (3)	Warning (4)
Info (4)	Info (6)

In the following example, Fabric OS messages map to local7 facility level 7 in the `/etc/syslog.conf` file:

```
local7.emerg      /var/adm/swcritical
local7.alert      /var/adm/alert7
local7.crit       /var/adm/crit7
local7.err        /var/adm/swerror
local7.warning    /var/adm/swwarning
local7.notice     /var/adm/notice7
local7.info       /var/adm/swinfo
local7.debug      /var/adm/debug7
```

If you prefer to map Fabric OS severities to a different UNIX local7 facility level, see [“Setting the facility level”](#) on page 139.

Configuring the switch

Configuring the switch involves specifying `syslogd` hosts and, optionally, setting the facility level. You can also remove a host from the list of `syslogd` hosts.

Specifying `syslogd` hosts

1. Connect to the switch and log in as admin.
2. Issue the `syslogDipAdd` command and specify an IP address.
3. Verify that the IP address was entered correctly, using the `syslogDipShow` command.

You can specify up to six host IP addresses for storing syslog messages, as shown in the following example:

```
switch:admin> syslogdipadd 10.1.2.1
switch:admin> syslogdipadd 10.1.2.2
switch:admin> syslogdipadd 10.1.2.3
switch:admin> syslogdipadd 10.1.2.4
switch:admin> syslogdipadd 10.1.2.5
switch:admin> syslogdipadd 10.1.2.6
switch:admin> syslogdipshow
syslog.IP.address.1 10.1.2.1
syslog.IP.address.2 10.1.2.2
syslog.IP.address.3 10.1.2.3
syslog.IP.address.4 10.1.2.4
syslog.IP.address.5 10.1.2.5
syslog.IP.address.6 10.1.2.6
```

Setting the facility level

You must set the facility level only if you specified a facility other than local7 in the host `/etc/syslog.conf` file.

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
syslogdfacility -l n
```

where *n* is a number from 0 through 7, indicating a UNIX local7 facility. The default is 7.

Removing a `syslogd` host from the list

1. Connect to the switch and log in as admin.
2. Issue the `syslogDipRemove` command:

```
switch:admin> syslogdipremove 10.1.2.1
```

3. Verify the IP address was deleted by issuing the `syslogDipShow` command.

Viewing and saving diagnostic information

Issue the `supportShow` command to dump important diagnostic and status information to the session screen, where you can review it or capture its data.

To save a set of files that customer support technicians can use to further diagnose the switch condition, issue the `supportSave` command. The command prompts for an FTP server, packages the following files, and sends them to the specified server:

- The output of the `supportShow` command
- The contents of any trace dump files on the switch
- System message logs (from both the CP blades for HP StorageWorks SAN Directors)

See also “[Setting up automatic trace dump transfers](#)” on page 140.

Setting up automatic trace dump transfers

You can set up a switch so that diagnostic information is transferred to a remote server. If a problem then occurs, you can provide your customer support representative with the most detailed information possible. To ensure the best service, set up for automatic transfer as part of standard switch configuration, before a problem occurs.

Setting up for automatic transfer of diagnostic files involves the following tasks:

- Specify a remote server on which to store the files.
- Enable the automatic transfer of trace dumps to the server. (Trace dumps overwrite each other by default; sending them to a server preserves information that would otherwise be lost.)
- Set up a periodic checking of the remote server, so you are alerted if the server becomes unavailable and you can correct the problem.

After the setup is complete, you can run the `supportSave -c` command to save diagnostic information to the server (without specifying server details).

The following procedures describe in detail the tasks for setting up automatic transfer. For details on the commands, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

Specifying a remote server

1. Verify that the FTP service is running on the remote server.
2. Connect to the switch and log in as admin.
3. Issue the following command:

```
supportftp -s
```

4. Respond to the prompts as follows:

- For `Host Name`, enter the name or IP address of the server where the files are to be stored; for example, `192.1.2.3`.
- For `user name`, enter the user name of your account on the server; for example, `JohnDoe`.
- For `password`, enter your account password for the server.
- For `remote directory`, specify a path name for the remote directory. Absolute path names can be specified using forward slash (/). Relative path names create the directory in the user's home directory on UNIX servers, and they are created in the directory where the FTP server is running on Windows servers.

Enabling the automatic transfer of trace dumps

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
traceftp -e
```

Setting up periodic checking of the remote server

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
supportftp -t interval
```

where the *interval* is in hours. The minimum interval is 1 hour. Specify 0 hours to disable the checking feature.

Saving a comprehensive set of diagnostic files to the server

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
supportsave -c
```


10 Troubleshooting

Troubleshooting should begin at the center of the SAN—the fabric. Because switches are located between the hosts and storage devices and have visibility into both sides of the storage network, starting with them can help narrow the search path. After eliminating the possibility of a fault within the fabric, determine whether the problem is on the storage side or the host side, and continue a more detailed diagnosis from there. Using this approach can quickly pinpoint and isolate problems.

For example, if a host cannot detect a storage device, run a switch command (such as `switchShow`) to find out whether the storage device is logically connected to the switch. If not, focus first on the switch directly connecting to the storage device. Use your vendor-supplied storage diagnostic tools to better understand why the storage device is not visible to the switch. If the storage device can be detected by the switch, and the host still cannot detect the storage device, there is still a problem between the host and switch.

This chapter provides information on troubleshooting and the most common procedures used to diagnose and repair issues. It also includes specific troubleshooting scenarios as examples.

Most common problem areas

See [Table 23](#) for a list of the most common problem areas that arise within SANs and a list of tools that can be used to resolve them.

Table 23 Common troubleshooting problems and tools

Problem area	Items to investigate	Tools
Fabric	<ul style="list-style-type: none">• Missing devices• Marginal links (unstable connections)• Incorrect zoning configurations• Incorrect switch configurations	<ul style="list-style-type: none">• Switch LEDs• Switch commands (for example, <code>switchShow</code> or <code>nsAllShow</code>) for diagnostics• Web or GUI-based monitoring and management software tools
Storage Devices	<ul style="list-style-type: none">• Physical issues between switch and devices• Incorrect storage software configurations	<ul style="list-style-type: none">• Device LEDs• Storage diagnostic tools• Switch commands (for example, <code>switchShow</code> or <code>nsAllShow</code>) for diagnostics
Hosts	<ul style="list-style-type: none">• Incorrect host bus adapter installation• Incorrect device driver installation• Incorrect device driver configuration	<ul style="list-style-type: none">• Host adaptor LEDs• Host operating system diagnostic tools• Device driver diagnostic tools• Switch commands (for example, <code>switchShow</code> or <code>nsAllShow</code>) for diagnostics
Storage management applications	Incorrect installation and configuration of the storage devices that the software references. For example, if using a volume-management application, check for: <ul style="list-style-type: none">• Incorrect volume installation• Incorrect volume configuration	Application-specific tools and resources

Gathering information for technical support

If you are troubleshooting a production system, you need to gather data quickly. As soon as a problem is observed, perform the following tasks (if you are using a dual-CP system, run the commands on both CPs):

1. Issue the `supportSave` command to save RASLOG, TRACE, and `supportShow` (active CP only) information for the local CP to a remote FTP location.

On a dual-CP system, only the local CP information is saved and `supportShow` information is not available on the active CP.

For details about the `supportSave` command, see “[Viewing and saving diagnostic information](#)” on page 140 and “[Setting up automatic trace dump transfers](#)” on page 140.

2. Issue the `pdShow` command to display data from a panic dump file.

The panic dump file contains information that might be useful to determine the cause of the system panic.

3. Issue the `saveCore` command to save or remove core files created by daemons.

For details about these commands, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

Troubleshooting questions

Common steps to take and questions to ask when troubleshooting a system problem are the following:

1. Determine the current Fabric OS level.
2. Determine the switch hardware version.
3. Determine whether the switch is operational.
4. Answer the following impact assessment and urgency questions:
 - Is the switch down?
 - Is it a standalone switch?
 - How large is the fabric?
 - Is the fabric redundant?

5. Run the `supportSave` command.

See “[Viewing and saving diagnostic information](#)” on page 140 and “[Setting up automatic trace dump transfers](#)” on page 140.

6. Document the sequence of events by answering the following questions:
 - What happened just prior to the problem?
 - Is the problem reproducible?
 - If so, what are the steps to produce the problem?
 - What configuration was in place when the problem occurred?
7. Determine whether a failover occurred.
8. Determine whether security was enabled.
9. Determine whether POST was enabled.
10. Determine whether serial port (console) logs were available.
11. Determine which CP blade was active. (This is applicable only to the Core Switch 2/64 and SAN Director 2/128.)
12. Determine what and when the last actions or changes were made to the system environment.

Use the following steps to retrieve as much of the informational items as possible before contacting the SAN technical support vendor.

1. Determine the following switch information:
 - Serial number (located on the chassis)
 - WWN (obtain using `licenseIdShow` or `wwn` commands)
 - Fabric OS version (obtain using `version` command)
 - Switch configuration settings
 - Command `supportSave` output

- Commands `pdShow` and `save` Core output
- 2. Determine the following host information:
 - OS version and patch level
 - HBA type
 - HBA firmware version
 - HBA driver version
 - Configuration settings
- 3. Determine the following storage information:
 - Disk/tape type
 - Disk/tape firmware level
 - Controller type
 - Controller firmware level
 - Configuration settings
 - Storage software (such as EMC Control Center, Veritas SPC, and the like)

Analyzing connection problems

If a host is unable to detect its target (for example, a storage or tape device), begin troubleshooting the problem in the middle of the data path. Determine whether the problem is above or below the starting point, and then continue to divide the suspected problem path in half until you can pinpoint the problem.

Use the following procedures to analyze the problem:

Checking the logical connection

1. Issue the `switchShow` command.
2. Review the output and determine whether the device is logically connected to the switch:
 - A device that is logically connected to the switch is registered as an F_Port or L_Port.
 - A device that is not logically connected to the switch is registered as something other than an Nx_Port.
3. If the missing device is logically connected, proceed to the next troubleshooting procedure ["Checking the Simple Name Server \(SNS\)"](#) on page 147.
4. If the missing device is not logically connected, check the device and everything on that side of the data path. Also see ["Correcting link failures"](#) on page 157.

Check all aspects of the host OS, the driver settings and binaries, the device Basic Input/Output System (BIOS) settings, the SFP, the cable going from the switch to the device, the SFP on the switch side of that cable, and all switch settings related to the device. See ["Checking for a loop initialization failure"](#) on page 158 as the next potential trouble spot.

Checking for Fibre Channel connectivity problems

Issue the `fcPing` command, which:

- Checks the zoning configuration for the two ports specified.
- Generates an ELS frame ECHO request to the source port specified and validates the response.
- Generates an ELS ECHO request to the destination port specified and validates the response.

Regardless of the device's zoning, the `fcPing` command sends the ELS frame to the destination port. A device can take any one of the following actions:

- Send an ELS Accept to the ELS request.
- Send an ELS Reject to the ELS request.
- Ignore the ELS request.

There are some devices that do not support the ELS ECHO request. In these cases, the device either does not respond to the request or it sends an ELS reject. When a device does not respond to the ELS request, further debugging is required; however, do not assume that the device is not connected to the Fibre Channel.

The following is sample output from the `fcPing` command in which one device accepts the request and another device rejects the request:

```
switch:admin> fcping 10:00:00:00:c9:29:0e:c4 21:00:00:20:37:25:ad:05
Source:      10:00:00:00:c9:29:0e:c4
Destination: 21:00:00:20:37:25:ad:05
Zone Check:  Not Zoned

Pinging 10:00:00:00:c9:29:0e:c4 [0x20800] with 12 bytes of data:
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1162 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1013 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1442 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1052 usec
received reply from 10:00:00:00:c9:29:0e:c4: 12 bytes time:1012 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1012/1136/1442 usec

Pinging 21:00:00:20:37:25:ad:05 [0x211e8] with 12 bytes of data:
Request rejected
Request rejected
Request rejected
Request rejected
Request rejected
5 frames sent, 0 frames received, 5 frames rejected, 0 frames timeout
Round-trip min/avg/max = 0/0/0 usec
switch:admin>
```

The following is sample output from the `fcPing` command in which one device accepts the request and another device does not respond to the request:

```
switch:admin> fcping 0x020800 22:00:00:04:cf:75:63:85
Source:      0x020800
Destination: 22:00:00:04:cf:75:63:85
Zone Check:  Zoned

Pinging 0x020800 with 12 bytes of data:
received reply from 0x020800: 12 bytes time:1159 usec
received reply from 0x020800: 12 bytes time:1006 usec
received reply from 0x020800: 12 bytes time:1008 usec
received reply from 0x020800: 12 bytes time:1038 usec
received reply from 0x020800: 12 bytes time:1010 usec
5 frames sent, 5 frames received, 0 frames rejected, 0 frames timeout
Round-trip min/avg/max = 1006/1044/1159 usec

Pinging 22:00:00:04:cf:75:63:85 [0x217d9] with 12 bytes of data:
Request timed out
Request timed out
Request timed out
Request timed out
Request timed out
5 frames sent, 0 frames received, 0 frames rejected, 5 frames timeout
Round-trip min/avg/max = 0/0/0 usec
switch:admin>
```

For details about the `fcPing` command, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

Checking the Simple Name Server (SNS)

1. Issue the `nsShow` command on the switch to which the device is attached.

For example:

```
The Local Name Server has 9 entries {

  Type Pid   COS      PortName          NodeName          TTL(sec)

*N  021a00;   2,3;20:00:00:e0:69:f0:07:c6;10:00:00:e0:69:f0:07:c6; 895
    Fabric Port Name: 20:0a:00:60:69:10:8d:fd
NL  051edc;   3;21:00:00:20:37:d9:77:96;20:00:00:20:37:d9:77:96; na
    FC4s: FCP [SEAGATE ST318304FC 0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee0;   3;21:00:00:20:37:d9:73:0f;20:00:00:20:37:d9:73:0f; na
    FC4s: FCP [SEAGATE ST318304FC 0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee1;   3;21:00:00:20:37:d9:76:b3;20:00:00:20:37:d9:76:b3; na
    FC4s: FCP [SEAGATE ST318304FC 0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee2;   3;21:00:00:20:37:d9:77:5a;20:00:00:20:37:d9:77:5a; na
    FC4s: FCP [SEAGATE ST318304FC 0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee4;   3;21:00:00:20:37:d9:74:d7;20:00:00:20:37:d9:74:d7; na
    FC4s: FCP [SEAGATE ST318304FC 0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051ee8;   3;21:00:00:20:37:d9:6f:eb;20:00:00:20:37:d9:6f:eb; na
    FC4s: FCP [SEAGATE ST318304FC 0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
NL  051eef;   3;21:00:00:20:37:d9:77:45;20:00:00:20:37:d9:77:45; na
    FC4s: FCP [SEAGATE ST318304FC 0005]

    Fabric Port Name: 20:0e:00:60:69:10:9b:5b
N   051f00;   2,3;50:06:04:82:bc:01:9a:0c;50:06:04:82:bc:01:9a:0c; na
    FC4s: FCP [EMC SYMMETRIX 5267]

    Fabric Port Name: 20:0f:00:60:69:10:9b:5b
```

2. Look for the device in the SNS list, which lists the nodes connected to that switch, allowing you to determine whether a particular node is accessible on the network.
 - If the device is not present in the SNS list, the problem is between the storage device and the switch. There might be a timeout communication problem between edge devices and the SNS, or there might be a login issue. First check the edge device documentation to determine whether there is a timeout setting or parameter that can be reconfigured. Check the port log for NS registration information and FCP probing failures (using the `fcPProbeShow` command). If these queries do not help solve the problem, contact the support organization for the product that appears to be inaccessible.
 - If the device is listed in the SNS list, the problem is between the storage device and the host. There might be a zoning mismatch or a host/storage issue. Proceed to ["Checking for zoning problems"](#) on page 148.
3. Issue the `portLoginShow` command to check the port login status.
4. Issue the `fcPProbeShow` command to display the FCP probing information for the devices attached to the specified F_Port or FL_Port.

This information includes the number of successful logins and SCSI INQUIRY commands sent over this port and a list of the attached devices.
5. Check the port log to determine whether the device sent the Fabric Login (FLOGI) frame to the switch, and the switch probed the device.

Checking for zoning problems

1. Issue the `cfgActvShow` command to determine whether zoning is enabled.
If zoning is enabled, it is possible that the problem is being caused by zoning enforcement (for example, two devices in different zones cannot see each other).
2. Confirm that the specific edge devices that need to communicate with each other are in the same zone.
 - If they are in the same zone, perform the following tasks:
 - a. Issue the `portCamShow` command on the host port to verify that the target is present.
 - b. Issue the `portCamShow` command on the target.
 - c. Issue the `nsZoneMember` command on the host and target to determine whether the Name Server is aware that these devices are zoned together.
 - If they are not in the same zone and zoning is enabled, proceed to [step 3](#).
3. Resolve zoning conflicts by putting the devices into the same zoning configuration.
See ["Correcting zoning setup issues"](#) on page 149 for additional information.

Restoring a segmented fabric

Fabric segmentation is generally caused by:

- Incompatible fabric parameters (see ["Reconciling fabric parameters individually"](#) on page 148).
- Incorrect PID setting (see ["Configuring the PID format"](#) on page 213).
- Incompatible zoning configuration (see ["Checking for zoning problems"](#) on page 148).
- Domain ID conflict (see ["Reconciling a domain ID conflict"](#) on page 149).
- A switch in a secure fabric not running Secure Fabric OS. (See the *HP StorageWorks Secure Fabric OS administrator guide* for additional information.)

There are a number of settings that control the overall behavior and operation of the fabric. Some of these values, such as the domain ID, are assigned by the fabric and can differ from one switch to another in the fabric. Other parameters, such as the BB credit, can be changed for specific applications or operating environments, but must be the same among all switches to allow the formation of a fabric.

The following fabric parameters must be identical for a fabric to merge:

- R_A_TOV
- E_D_TOV
- Data field size
- Sequence level switching
- Disabled device probing
- Suppressed class F traffic
- Per-frame route priority
- Long-distance fabric (not necessary on Bloom-based fabrics)
- BB credit
- PID format

Reconciling fabric parameters individually

1. Log in to one of the segmented switches as admin (switch A).
2. Issue the `configShow` command.
3. Log in to another switch (switch B) in the same fabric as admin.
4. Issue the `configShow` command.
5. Compare the two switch configurations line by line and look for differences. Do this by comparing the two telnet windows or by printing the `configShow` output. Verify that the fabric parameter settings (see ["Restoring a segmented fabric"](#)) are the same for both switches.
6. Connect to the segmented switch after the discrepancy is identified.
7. Disable the switch by issuing the `switchDisable` command.

8. Issue the `configure` command to edit the fabric parameters for the segmented switch.
See the *HP StorageWorks Fabric OS 5.x command reference guide* for detailed information.
9. Enable the switch by issuing the `switchEnable` command.

You can also reconcile fabric parameters by issuing the `configUpload` command for each switch.

Downloading a correct configuration

You can restore a segmented fabric by downloading a previously saved correct backup configuration to the switch. Downloading in this manner reconciles any discrepancy in the fabric parameters and allows the segmented switch to rejoin the main fabric. For details on uploading and downloading configurations, see *"Maintaining configurations"* on page 73.

Reconciling a domain ID conflict

If a domain ID conflict appears, the conflict is reported only at the point where the two fabrics are physically connected. However, there might be several conflicting domain IDs, which appears as soon as the initial conflict is resolved.

Typically, the fabric resolves domain conflicts during fabric merges or builds unless IDID is configured. If IDID is enabled, switches that cannot be programmed with a unique domain ID are segmented out. Check each switch that has IDID enabled and make sure their domain IDs are unique within the configuration.

Repeat this procedure until all domain ID conflicts are resolved:

1. Issue the `fabricShow` command on a switch from one of the fabrics.
2. In a separate telnet window, issue the `fabricShow` command on a switch from the second fabric.
3. Compare the `fabricShow` output from the two fabrics. Note the number of domain ID conflicts; there might be several duplicate domain IDs that need to be changed.
4. Determine which switches have domain overlap and change the domain IDs for each of those switches.
5. Chose the fabric on which to change the duplicate domain ID; connect to the conflicting switch in that fabric.
6. Issue the `switchDisable` command.
7. Issue the `switchEnable` command.
This enables the joining switch to obtain a new domain ID as part of the process of coming online. The fabric principal switch allocates the next available domain ID to the new switch during this process.
8. Repeat [step 5](#) through [step 7](#) if additional switches have conflicting domain IDs.

Correcting zoning setup issues

The types of zone configuration discrepancies that can cause segmentation are listed in [Table 24](#).

Table 24 Types of zone discrepancies

Conflict	Description
Configuration mismatch	Occurs when zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
Type mismatch	Occurs when the name of a zone object in one fabric is also used for a different type of zone object in the other fabric. A zone object is any device in a zone.
Content mismatch	Occurs when the definition in one fabric is different from the definition of a zone object with the same name in the other fabric.

[Table 25](#) summarizes commands that are useful for debugging zoning issues.

Table 25 Commands for debugging zoning

Command	Function
aliCreate	Use to create a zone alias.
aliDelete	Use to delete a zone alias.
cfgCreate	Use to create a zone configuration.
cfgShow	Displays zoning configuration.
licenseShow	Displays current license keys and associated (licensed) products.
switchShow	Displays currently enabled configuration and any E_Port segmentations due to zone conflicts.
zoneAdd	Use to add a member to an existing zone.
zoneCreate	Use to create a zone. Before a zone becomes active, the <code>cfgSave</code> and <code>cfgEnable</code> commands must be used.
zoneHelp	Displays help information for zone commands.
zoneShow	Displays zone information.

See “[Administering advanced zoning](#)” on page 177 for additional information about setting up zoning properly. Also see the *HP StorageWorks Fabric OS 5.x command reference guide* for details about zoning commands.

Correcting a fabric merge problem quickly

You can correct zone conflicts by using the `cfgClear` command to clear the zoning database.

To correct a merge conflict without disrupting the fabric, first verify fabric merge problem, and then edit zone configuration members, and then reorder the zone member list.

△ **CAUTION:** This is a disruptive procedure.

1. Determine which switches have the incorrect zoning configuration, and then log in to the switches as admin.
 2. Issue the `switchDisable` command.
 3. Issue the `cfgDisable` command.
 4. Issue the `cfgClear` command.
-

△ **CAUTION:** This command clears the zoning database on the affected switches.

5. Issue the `switchEnable` command.
This forces a zone merge and populates the switches with the desired zoning database. The two fabrics merge again.

Verifying a fabric merge problem

1. Issue the `switchShow` command to validate that the segmentation is due to a zone issue.
2. See [Table 24](#) on page 149 to view the different types of zone discrepancies.

Editing zone configuration members

1. Log in to one of the switches in a segmented fabric as admin.
2. Issue the `cfgShow` command and print its output.
3. Start another telnet session and connect to the next fabric as an administrator.
4. Issue the `cfgShow` command and print its output.
5. Compare the two fabric zone configurations line by line and look for an incompatible configuration.
6. Connect to one of the fabrics.
7. Issue zone configure edit commands to edit the fabric zone configuration for the segmented switch (see [Table 25](#) on page 150 for specific commands).

If the zoneset members between two switches are not listed in the same order in both configurations, the configurations are considered a mismatch; this results in the switches being segmented in the fabric.

For example:

`[cfg1 = z1; z2]` is different from `[cfg1 = z2; z1]`, even though the members of the configuration are the same.

One simple approach to making sure that the zoneset members are in the same order is to keep the members in alphabetical order.

Reordering the zone member list

1. Use the output from the `cfgShow` for both switches by comparing the order that the zone members are listed.
Members must be listed in the same order.
2. Rearrange zone members so that the configuration for both switches is the same. Arrange zone members in alphabetical order, if possible.

Recognizing MQ errors

Identify a message queue (MQ) error message by looking for the two letters M and Q in the error message.

For example:

```
2004/08/24-10:04:42, [MQ-1004], 218,, ERROR, ras007, mqRead, queue =  
raslog-test-string0123456-raslog, queue I  
D = 1, type = 2
```

MQ errors can result in devices dropping from the SNS or can prevent a switch from joining the fabric. MQ errors are rare and difficult to troubleshoot; resolve them by working with the switch supplier. When MQ errors are encountered, issue the `supportSave` command to capture debug information about the switch, and then forward the `supportSave` data to the switch supplier for further investigation.

Correcting I²C bus errors

I²C bus errors indicate defective hardware; the specific item is listed in the error message. See the *HP StorageWorks Fabric OS 5.x diagnostics and system error messages reference guide* for information specific to the error that was received. Some CPT and Environmental Monitor (EM) messages contain I²C-related information.

If the I²C message does not indicate the specific hardware that might be failing, begin debugging the hardware, as this is the most likely cause. The next sections provide procedures for debugging the hardware.

Checking fan components

1. Log in to the switch as user.
2. Issue the `fanShow` command.
3. Check the fan status and speed output.

If any of the fan speeds display abnormal RPMs, replace the fan FRU.

Checking the switch temperature

1. Log in to the switch as user.
2. Issue the `tempShow` command.
3. Check the temperature output. Look for indications of high or low temperatures.

Checking the power supply

1. Log in to the switch as user.
2. Issue the `psShow` command.
3. Check the power supply status.

See the appropriate SAN Switch installation guide for details regarding the power supply status.

If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible.

Checking the temperature, fan, and power supply

1. Log in to the switch as user.
2. Issue the `sensorShow` command.

See the *HP StorageWorks Fabric OS 5.x command reference guide* for details regarding the sensor numbers.

3. Check the temperature output. Look for indications of high or low temperatures.
4. Check the fan speed output.

If any of the fan speeds display abnormal RPMs, replace the fan FRU.

5. Check the power supply status.

If any of the power supplies show a status other than OK, consider replacing the power supply as soon as possible.

Correcting device login issues

To try to pinpoint problems with device logins, use the following procedure:

1. Log in to the switch as admin.
2. Issue the `switchShow` command, and then check for correct logins.

For example:

```
switch:admin> switchshow
switchName:      sw094135
switchType:      26.1
switchState:      Online
switchMode:      Native
switchRole:      Principal
switchDomain:     126
switchId:        fffc7e
switchWwn:       10:00:00:05:1e:34:00:69
zoning:          ON (cfg_em)
switchBeacon:    OFF
Port   Media Speed State
=====
  0    id    N1   Online   E-Port  10:00:00:60:69:11:f9:fc "2800_116"
  1    id    1G   Online   E-Port  10:00:00:60:69:11:f9:fc "2800_116"
  2    id    N2   No_Light
  3    id    2G   No_Light
  4    id    N2   Online   E-Port  (Trunk port, master is Port  5)
  5    id    N2   Online   E-Port  10:00:00:05:1e:34:00:8b "Dazz125"
(downstream)(Trunk master)
  6    id    N2   No_Light
  7    id    N2   No_Light
  8    id    N1   Online   L-Port  4 public, 1 private, 1 phantom
  9    id    N2   No_Light
 10    id    N2   Online   G-Port
 11    id    N2   Online   F-Port  10:00:00:01:c9:28:c7:01
 12    id    N1   Online   L-Port  4 public, 1 private, 1 phantom
 13    --    N2   No_Module
 14    id    N2   Online   E-Port  (Trunk port, master is Port 15)
 15    id    N2   Online   E-Port  10:00:00:60:69:90:03:17 "TERM_113"
(downstream)(Trunk master)
```

3. Issue the portCfgShow command to see how the port is configured.

For example:

```
sw094135:admin> portcfgshow
Ports of Slot 0    0  1  2  3    4  5  6  7    8  9 10 11    12 13 14 15
-----+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
Speed          AN 1G AN 2G   AN AN AN AN   AN AN AN AN   AN AN AN AN
Trunk Port      ON ON .. ON   ON ON ON ON   ON ON ON ON   ON ON ON ON
Long Distance   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
VC Link Init    .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked L_Port   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked G_Port   .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Disabled E_Port .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
ISL R_RDY Mode  .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Persistent Disable.. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..
Locked Loop HD  .. .. .. ..   .. .. .. ..   .. .. .. ..   .. .. .. ..

                                where AN:AutoNegotiate, ..:OFF, ?:INVALID.
p                                LM:L0.5
```

4. Issue the portErrShow command and then check for errors that can cause login problems.

For example:

```
sw094135:admin> porterrshow
```

	frames		enc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy
	tx	rx	in	err	shrt	long	eof	out	c3	fail	sync	sig		
0:	38	75	0	0	0	0	0	0	0	9	11	0	0	0
1:	110	73	0	0	0	0	0	0	0	9	11	0	0	0
2:	0	0	0	0	0	0	0	38	0	4	0	2	0	0
3:	0	0	0	0	0	0	0	0	0	4	1	2	0	0
4:	59m	102	0	0	0	0	0	0	0	4	0	0	0	0
5:	59m	103	0	0	0	0	0	0	0	3	0	0	0	0
6:	0	0	0	0	0	0	0	21	0	3	0	0	0	0
7:	0	0	0	0	0	0	0	58	0	3	0	0	0	0
8:	81	19k	0	0	0	0	0	3.0m	0	5	43	0	0	0
9:	0	0	0	0	0	0	0	29	0	3	0	0	0	0
10:	12m	68m	0	0	0	0	0	13	43m	8	1	1	0	0
11:	30m	33m	0	0	0	0	0	0	0	8	1	1	0	0
12:	89	25k	0	0	0	0	0	2.9m	0	7	43	0	0	0
13:	0	0	0	0	0	0	0	0	0	3	0	0	0	0
14:	29m	82m	0	0	0	0	0	0	1.2m	4	1	1	0	0
15:	29m	81m	0	0	0	0	0	0	1.1m	4	1	1	0	0

- A high number of errors relative to the frames transmitted and frames received can indicate a marginal link (see ["Correcting marginal links"](#) on page 159 for additional information).
- A steadily increasing number of errors can indicate a problem. Track errors by sampling the port errors every five or ten minutes.

5. Issue the `portFlagsShow` command and then check to see how a port has logged in and where a login failed (if a failure occurred).

For example:


```
sw094135:admin> portflagsshow
```

Port	SNMP	Physical	Flags
0:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED			
1:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED			
2:	Offline	No_Light	PRESENT U_PORT LED
3:	Offline	No_Light	PRESENT U_PORT LED
4:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED			
5:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED			
6:	Offline	No_Light	PRESENT U_PORT LED
7:	Offline	No_Light	PRESENT U_PORT LED
8:	Online	In_Sync	PRESENT ACTIVE F_PORT L_PORT U_PORT LOGICAL_ONLINE
LOGIN NOELP LED ACCEPT			
9:	Offline	No_Light	PRESENT U_PORT LED
10:	Online	In_Sync	PRESENT ACTIVE G_PORT U_PORT LOGIN LED
11:	Online	In_Sync	PRESENT ACTIVE F_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN NOELP LED ACCEPT			
12:	Online	In_Sync	PRESENT ACTIVE F_PORT L_PORT U_PORT LOGICAL_ONLINE
LOGIN NOELP LED ACCEPT			
13:	Offline	No_Module	PRESENT U_PORT LED
14:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED			
15:	Online	In_Sync	PRESENT ACTIVE E_PORT G_PORT U_PORT LOGICAL_ONLINE
LOGIN LED			

6. Issue the `portLogDumpPort portid` command, where `portID` is the port number, and then view the device to switch communication.

For example:

```
sw094135:admin> portlogdumpport 10
time          task      event  port cmd  args
-----
12:38:21.590  SPEE      sn      10   WS   00000000,00000000,00000000
12:38:21.591  SPEE      sn      10   WS   000000ee,00000000,00000000
12:38:21.611  SPEE      sn      10   WS   00000001,00000000,00000000
12:38:21.871  SPEE      sn      10   NC   00000002,00000000,00000001
12:38:21.872  LOOP     loopscn 10   LIP   8002
12:38:22.171  LOOP     loopscn 10   TMO   2
12:38:22.171  INTR     pstate 10   LF2
12:38:22.172  INTR     pstate 10   OL2
12:38:22.172  INTR     pstate 10   LR3
12:38:22.172  INTR     pstate 10   AC
12:38:22.172  PORT     scn      10   11   00000000,00000000,00000002
12:38:22.311  PORT     scn      10   1    00000000,00000000,00000001
12:38:22.311  PORT     debug    10   00000001,00654320,00000001,00000000
12:38:22.311  PORT     debug    10   00000001,00654320,00000002,00000000
12:38:22.311  PORT     debug    10   00000001,00654320,00000003,00000000
12:38:22.313  PORT     Tx       10   164  02ffffff,00ffffff,025effff,10000000
12:38:22.314  PORT     debug    10   00000001,00654320,00000003,00000000  *
7
12:38:28.312  PORT     Tx       10   164  02ffffff,00ffffff,028fffff,10000000
12:38:34.312  PORT     Tx       10   164  02ffffff,00ffffff,0293ffff,10000000
12:38:40.312  PORT     Tx       10   164  02ffffff,00ffffff,0299ffff,10000000
12:38:46.312  PORT     Tx       10   164  02ffffff,00ffffff,029bffff,10000000
12:38:52.312  PORT     Tx       10   164  02ffffff,00ffffff,029dffff,10000000
12:38:58.312  PORT     Tx       10   164  02ffffff,00ffffff,02acffff,10000000
12:39:04.322  INTR     pstate 10   LR1
12:39:04.323  INTR     pstate 10   LR3
12:39:04.323  INTR     pstate 10   AC
12:39:04.323  PORT     scn      10   11   00000000,00000000,00000002
sw094135:root>
```

 **NOTE:** See “[Viewing the port log](#)” on page 137 for overview information about the output of the portLogDump command.

Identifying media-related issues

This section provides procedures that help pinpoint any media-related issues in the fabric. The tests listed in [Table 26](#) are a combination of structural and functional tests that can be used to provide an overview of the hardware components and help identify media-related issues.

- Structural tests perform basic testing of the switch circuit. If a structural test fails, replace the main board or port blade.
- Functional tests verify the intended operational behavior of the switch by virtue of running frames through ports or bypass circuitry.

Table 26 Component test descriptions

Test name	Operands	Checks
crossporttest	[-nframes <i>count</i>] [-lb_mode <i>mode</i>][-spd_mode <i>mode</i>] [-gbic_mode <i>mode</i>] [-norestore <i>mode</i>] [-ports <i>itemlist</i>]	Functional test of port external transmit and receive path. The crossport is set to loopback using an external cable by default. However, this command can be used to check internal components by setting the <i>lb</i> operand to 5.
fporttest	[-nframes <i>count</i>] [-ports <i>itemlist</i>] [-seed <i>payload_pattern</i>] [-width <i>pattern_width</i>] [-size <i>pattern_size</i>]	Tests component to and from an HBA. Used to test online F_Port devices, N_Port devices, SFPs, and GBICs.
loopporttest	[-nframes <i>count</i>] [-ports <i>itemlist</i>][-seed <i>payload_pattern</i>] [-width <i>pattern_width</i>]	Tests only components attached to switches that are on an FC-AL.
spinfab	[<i>nMillionFrames</i> [, <i>ePortBeg</i> [, <i>ePortEnd</i> [, <i>setFail</i>]]]]	Tests components between switches, such as ISLs, SFPs, and GBICs to and from a neighbor switch.

The following procedures are for checking switch-specific components.

Testing a port's external transmit and receive path

1. Connect to the switch and log in as admin.
2. Connect the port you want to test to any other switch port with the cable you want to test.
3. Issue the `crossPortTest` command with the following operands.)

This is a partial list. See the *HP StorageWorks Fabric OS 5.x command reference guide* for additional command information:

- [-nframes *count*], which specifies the number of frames to send
- [-lb_mode *mode*], which selects the loopback point for the test
- [-spd_mode *mode*], which selects the speed mode for the test
- [-ports *itemlist*], which specifies a list of user ports to test

For example:

```
switch:admin> crossporttest
Running Cross Port Test .... passed.
```

Testing a switch's internal components

1. Connect to the switch and log in as admin.
2. Connect the port you want to test to any other switch port with the cable you want to test.
3. Issue the `crossporttest -lb_mode 5` command where 5 is the operand that causes the test to be run on the internal switch components.

The following is partial list of operands—see the *HP StorageWorks Fabric OS 5.x command reference guide* for additional command information:

- [-nframes *count*], specifies the number of frames to send
- [-lb_mode *mode*], selects the loopback point for the test
- [-spd_mode *mode*], selects the speed mode for the test
- [-ports *itemlist*], specifies a list of user ports to test

Testing components to and from the HBA

1. Connect to the switch and log in as admin.
2. Issue the `fPortTest` command.

See the *HP StorageWorks Fabric OS 5.x command reference guide* for information on the command options. The following example executes the `fPortTest` command 100 times on port 8 with payload pattern 0xaa55, pattern width 2 (meaning word width) and a default payload size of 512 bytes:

```
switchname:admin> fporttest 100,8,0xaa55,2, 512
Will use pattern: aa55 aa55 aa55 aa55 aa55 aa55 ...
Running fPortTest .....
port 8 test passed.
value = 0
```

Table 27 provides a list of additional tests that can be used to determine the switch components that are not functioning properly. See the *HP StorageWorks Fabric OS 5.x command reference guide* for additional command information.

Table 27 Switch component tests

Test	Function
portloopbacktest	Functional test of port N to N path.
portregtest	Read and write test of the ASIC SRAMs and registers.
spinsilk	Functional test of internal and external transmit and receive paths at full speed.
sramretentiontest	Test to verify that the data written into the miscellaneous SRAMs in the ASIC are retained after a 10-second wait.
crossporttest	Test to verify that the functional components of the switch.
turboramtest	Test to verify that the on chip SRAM located in the 2 Gbit/sec ASIC is using the Turbo-Ram BIST circuitry. These same SRAMs are tested by <code>portregtest</code> and <code>sramretentiontest</code> using PCI operations, but for this test the BIST controller is able to perform the SRAM write and read operations at a much faster rate.
statstest	Test to verify that the ASIC statistics counter logic.
Related Switch Test Option:	
itemlist	Option to restrict the items to be tested to a smaller set of parameter values that you pass to the switch.


Correcting link failures

A link failure occurs when a server or storage device is connected to a switch, but the link between the server or storage device and the switch does not come up. This prevents the server or storage device from communicating through the switch.

If the `switchShow` command or LEDs indicate that the link has not come up properly, use one or more of the following procedures.

Determining whether the negotiation was successfully completed

The port negotiates the link speed with the opposite side. The negotiation usually completes in 1–2 seconds; however, sometimes the speed negotiation fails.

 **NOTE:** Skip this procedure if the port speed is set to a static speed through the `portCfgSpeed` command.

1. Issue the `portCfgShow` command to display the port speed settings of all the ports.
2. Issue the `switchShow` command to determine whether the port has module light.
3. Determine whether the port at 1 Gig/sec completes by issuing the `portCfgSpeed` command, and then change the port speed to 2 Gig/sec.
This should correct the negotiation by the port setting to one speed.
4. Issue the `portLogShow` or `portLogDump` command.
5. Check the events area of the output.

The first of the following examples is 1 Gbit/sec; the second example is 2 Gbit/sec:

```
14:38:51.976  SPEE sn <Port#>  NC  00000001,00000000,00000001
```

```
14:39:39.227  SPEE          sn      <Port#>  NC  00000002,00000000,00000001
```

where:

- `sn` indicates a speed negotiation.
- `NC` indicates negotiation completion.
- `01` or `02` indicate the speed that has been negotiated.

If these fields do not appear, proceed to the [step 6](#).

6. Correct the negotiation by issuing the `portCfgSpeed [slotnumber/]portnumber, speed_level` command if the fields in [step 5](#) do not appear.

Checking for a loop initialization failure

1. Verify the port is an L_Port:
 - a. Issue the `switchShow` command.
 - b. Check the comment field of the output to verify that the switch port indicates an L_Port.
If a loop device is connected to the switch, the switch port must be initialized as an L_Port.
 - c. Check to ensure that the port state is online; otherwise, check for link failures.
2. Verify the loop initialization if the port is not an L_port:
 - a. Issue the `portLogShow` or `portLogDump` command.
 - b. Check argument number four for the Loop Initialization Soft Assigned (LISA) frame (0x11050100).
For example:

```
termB:admin> portlogdumpport 4
time          task          event  port cmd  args
-----
11:40:02.078  PORT          Rx3     23   20
22000000,00000000,ffffffff,11050100 Received LISA frame
```

The LISA frame indicates that the loop initialization is complete.

3. Skip point-to-point initialization by issuing the `portCfgLport` command.
The switch changes to point-to-point initialization after the LISA phase of the loop initialization. This behavior sometimes causes trouble with old HBAs. If that is the case, skip point-to-point initialization.

Checking for a point-to-point initialization failure

1. Issue the `switchShow` command to confirm that the port is active and has a module that is synchronized.
If a fabric device or another switch is connected to the switch, the port must be online.
2. Issue the `portLogShow` or `portLogDump` command.

3. Verify that the event area for the port state entry is `pstate`.

The command entry `AC` indicates that the port has completed point-to-point initialization. For example:

```
termB:root> portlogdumpport 4
time          task          event  port cmd  args
-----
11:38:21.726  INTR          pstate  4    AC
```

4. Skip over the loop initialization phase.

After becoming an active port, the port becomes an `F_Port` or an `E_Port`, depending on the device on the opposite side. If the opposite device is a fabric device, the port becomes an `F_Port`. If the opposite device is another switch, the port becomes an `E_Port`.

Some fabric devices have problems with loop initialization. If this is evident, issue the `portCfgPort port_#, 1` command.

Correcting a port that came up in the wrong mode

1. Issue the `switchShow` command.
2. Note the comment fields (see [Table 28](#)) and follow the suggested actions.

Table 28 Switchshow output and suggested action

Output	Suggested action
Disabled	Check the output from the <code>switchShow</code> command to determine whether the switch is disabled. If the switch is disabled (for example, due to persistent disable or security reasons), attempt to resolve the issue, and then issue the <code>portEnable</code> command.
Bypassed	Check the output from the <code>switchShow</code> command to determine whether the port is testing.
Loopback	Check the output from the <code>switchShow</code> command to determine whether the port is testing.
E_Port	If the opposite side is not another switch, the link came up in a wrong mode. Check the output from the <code>portLogShow</code> or <code>PortLogDump</code> command and identify the link initialization stage to determine where the initialization procedure went wrong.
F_Port	If the opposite side of the link is a fabric device, the link came up in a wrong mode. Check the output from <code>portLogShow</code> or <code>PortLogDump</code> command.
G_Port	The port did not come up as an <code>E_Port</code> or <code>F_Port</code> . Check the output from <code>portLogShow</code> or <code>PortLogDump</code> command and identify the link initialization stage where the initialization procedure went wrong.
L_Port	If the opposite side is not a loop device, the link came up in a wrong mode. Check the output from <code>portLogShow</code> or <code>PortLogDump</code> command and identify the link initialization stage where the initialization procedure went wrong.

Correcting marginal links

A marginal link involves the connection between the switch and the edge device. Isolating the exact cause of a marginal link involves analyzing and testing many of the components that make up the link (including the switch port, the switch SFP, the cable, the edge device, and the edge device SFP).

To troubleshoot a marginal link:

1. Issue the `portErrShow` command.

For example:

```
switch:admin> porterrshow
      frames  enc  crc  too  too  bad  enc  disc  link  loss  loss  frjt  fbsy
      tx   rx   in  err  shrt long  eof  out   c3 fail sync sig
sig=====
0:   22   24   0   0   0   0   0  1.5m  0   7   3   0   0   0
1:   22   24   0   0   0   0   0  1.2m  0   7   3   0   0   0
2:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
3:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
4:  149m  99m   0   0   0   0   0  448   0   7   6   0   0   0
5:  149m  99m   0   0   0   0   0  395   0   7   6   0   0   0
6:  147m  99m   0   0   0   0   0  706   0   7   6   0   0   0
7:  150m  99m   0   0   0   0   0  160   0   7   5   0   0   0
8:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
9:    0    0   0   0   0   0   0    0   0   0   0   0   0   0
10:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
11:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
12:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
13:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
14:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
15:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
32:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
33:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
34:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
35:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
36:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
37:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
38:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
39:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
40:  99m  146m   0   0   0   0   0  666   0   6  796   7   0   0
41:  99m  149m   0   0   0   0   0  15k   0   2  303   4   0   0
42:  99m  152m   0   0   0   0   0  665   0   2  221   5   0   0
43:  99m  147m   0   0   0   0   0  16k   0   2  144   4   0   0
44:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
45:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
46:   0    0   0   0   0   0   0    0   0   0   0   2   0   0
47:   0    0   0   0   0   0   0    0   0   0   0   0   0   0
switch:admin>
```

2. Establish whether there is a relatively high number of errors (such as CRC errors or ENC_OUT errors), or if there is a steadily increasing number of errors to confirm a marginal link.
3. If you suspect a marginal link, isolate the areas by moving the suspected marginal port cable to a different port on the switch.
If the problem stops or goes away, the switch port or the SFP is marginal (proceed to [step 4](#))
If the problem does not stop or go away, proceed to [step 7](#).
4. Replace the SFP on the marginal port.
5. Verify that you have an adapter to run the loopback test for the SFP. If you do not have an adapter, run the `portLoopBack` test on the marginal port using the loopback mode `1b=5`. See the *HP StorageWorks Fabric OS 5.x command reference guide* for additional information. [Table 29](#) describes the loopback modes.

Table 29 Loopback modes

Mode	Description
1	Port Loopback (loopback plugs)
2	External (SERDES) loopback

Table 29 Loopback modes (continued)

Mode	Description
5	Internal (parallel) loopback (indicates no external equipment)
7	Backend bypass and port loopback
8	Backend bypass and SERDES loopback
9	Backend bypass and internal loopback

6. Check the results of the loopback test and proceed as follows:
 - If the loopback test failed, the port is bad. Replace the port blade.
 - If the loopback test did not fail, the SFP is bad.
7. Optional: To rule out cabling issues:
 - a. Insert a new cable into the suspected marginal port.
 - b. Issue the `portErrShow` command to determine whether a problem still exists.
 - If the `portErrShow` output displays a normal number of generated errors, the issue is solved.
 - If the `portErrShow` output still displays a high number of generated errors, follow the troubleshooting procedures for the host or storage device.

Inaccurate information in the system message log

In rare instances, events gathered by the track change feature can report inaccurate information to the system message log.

For example, consider a case where a user enters a correct user name and password, but the login is rejected because the maximum number of users was reached. The system message log reports this login as successful.

If the maximum number of switch users is reached, the switch still performs correctly in that it rejects the login of additional users (even if they enter a correct user name and password).

However, in this limited example, the Track Change feature reports this event inaccurately to the system message log and the login appears successful. This scenario occurs only when the maximum number of users is reached; otherwise, the login information displayed in the system message log should reflect reality.

For information regarding enabling and disabling TC, see [“Tracking and controlling switch changes”](#) on page 35.

Port initialization and FCP auto-discovery process

The steps in the port initialization process represent a protocol used to discover the type of connected device and establish the port type. The possible port types are as follows:

- **U_Port:** Universal FC port. This port type is the base Fibre Channel port type and all unidentified or uninitiated ports are listed as U_Ports.
- **FL_Port:** Fabric Loop port. This port connects both public and private loop devices.
- **G_Port:** Generic port. This port acts a transition port for non-loop fabric-capable devices (E_Port and F_Port).
- **E_Port:** Expansion port. This port type is assigned to ISL links.
- **F_Port:** Fabric port. This port is assigned to fabric capable devices.

The FCP auto-discovery process enables private storage devices that accept PRLI to communicate in a fabric.

If device probing is enabled, the embedded port logins (PLOGIs) and attempts a PRLI into the device to retrieve information to enter into the Name Server. This enables private devices that do not FLOGI but accept PRLI to be entered in the Name Server and receive full fabric citizenship. Private devices that

accept PRLI represent a majority of storage targets. Private hosts require the QuickLoop feature, which is not available in Fabric OS 4.0.0 or later.

A fabric-capable device implicitly registers information with the Name Server during a FLOGI. These devices typically register information with the Name Server before querying for a device list. The embedded port still performs a PLOGI and attempts PRLI with these devices.

You can view the Name Server table in Advanced Web Tools by clicking Name Server in the fabric toolbar. See the *HP StorageWorks Fabric OS 5.x Advanced Web Tools administrator guide* for more information.

11 Administering extended fabrics

This chapter contains procedures for using the HP Extended Fabrics licensed feature, which extends the distance that ISLs can reach. To use extended ISL modes, you must first purchase and install the Extended Fabrics license. For details on obtaining and installing licensed features, see ["Maintaining licensed features"](#) on page 26.

About extended link buffer allocation

As the distance between switches and the link speed increase, additional buffer-to-buffer credits are required to maintain maximum performance. The number of credits reserved for a port depends on the switch model and on the extended ISL mode for which it is configured.

SAN Switch 2/8V, SAN Switch 2/16V, and SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director (FC2-16 port blades)

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director using FC2-16 port blades, each port group contains four ports and uses a common pool of credits. Because the number of credits available for use within each port group is limited, configuring ports for extended links on these models might cause other ports to become disabled if there are not enough buffer credits available; for example:

- If two 2-Gb/second ports in a group are configured for L1 mode, each is allocated sufficient buffer-to-buffer credits to cause the other two ports in the group to become disabled.
- A port connected to a device that is in loopback mode might become disabled for lack of buffers if another port in that group is set to L2 mode.

See Chapter 5, ["Configuring Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director"](#) for details about port blade nomenclature.

Brocade 4Gb SAN Switch for HP p-Class BladeSystem, SAN Switch 4/32, and 4/256 SAN Director (FC4-16 and FC4-32 port blades)

For the SAN Switch 4/32 and 4/256 SAN Director using FC4-16 and FC4-32 port blades, buffer credits are used by all ports on chip. Buffer-limited port technology allows all ports to remain operational, even when extended links are in use.

For the Brocade 4Gb SAN Switch for HP p-Class BladeSystem, buffer credits are available to all ports on the chip.

A buffer-limited port can come online with fewer buffer credits allocated than its configuration specifies, allowing it to operate at a reduced bandwidth instead of being disabled for lack of buffers.

Buffer-limited operation is supported for the L0 and LD extended ISL modes only, and is persistent across reboots, switch disabling and enabling, and port disabling and enabling.

Fabric considerations

Consider these items that affect the fabric when you configure extended ISLs:

- The extended link configuration mode, L2, can reach 100 km at a speed of 2 Gb/sec between Fabric OS 4.x switches. However, it supports up to 60 km only if the link is established between Fabric OS 3.x and 4.x switches.
- The standard distance and LD ISL modes cannot be enabled at the same time.
- Balance the number of LD ISL connections and core-to-edge ISL connections within a switch. Configuring LD ISLs between core and edge switches is possible, but HP does not recommend it.

- Starting with Fabric OS 4.4.0, VC translation link initialization (an option of the `portCfgLongDistance` command) is enabled by default for LD links. For previous Fabric OS versions that support this option, it was disabled by default. To avoid inconsistency in the fabric, make sure that this value is enabled on both ends of the link. To connect to switches running Fabric OS versions earlier than 4.0.2 and 3.0.2c, make sure that VC translation link initialization is disabled because these versions do not support it.

Choosing an extended ISL mode

Table 30 lists the extended ISL modes for switches that have a Bloom ASIC. You can configure extended ISL modes with the `portCfgLongDistance` command when the Extended Fabrics license is activated.

Table 30 Extended ISL modes: switches with Bloom ASIC

Mode	Description	Buffer Allocation		Distance at 1 Gbps	Distance at 2 Gbps	Earliest Fabric OS release	Extended Fabrics License Required?
		1 Gbps	2 Gbps				
L0	Level 0 static mode; the default	5 (26) ²	5 (26)	10 km	5 km	All	No
LE	Level E static mode; supports links beyond 5 km	13	19	N/A	10 km	3.x, 4.x	No
L0.5	Level 0.5 static mode (designated LM when listed with the <code>portCfgshow</code> command)	19	34	25 km	25 km	3.1.0, 4.1.0, 4.x, 5.x	Yes
L1	Level 1 static mode	27	54	50 km	50 km	All	Yes
L2	Level 2 static mode	60	64	100 km	60 km	All	Yes
LD ¹	Dynamic mode uses automatic distance detection for a user-specified distance	Auto	Auto	Auto	Auto	3.1.0, 4.1.0, 4.4.0, 5.x (depending on the model)	Yes

¹The dynamic long-distance mode (LD) configures the number of buffer credits required, based on the actual link distance.

²For each data channel (in this case there are 4) there are 5 credits, plus 6 extra credits.

Table 31 lists the extended ISL modes for switches that have a Goldeneye ASIC.

Table 31 Extended ISL modes: switches with Goldeneye ASIC

Mode	Buffer Allocation			Distance at 1 Gbps	Distance at 2 Gbps	Distance at 4 Gbps	Earliest Fabric OS release	Extended fabrics license required?
	1 Gbps	2 Gbps	4 Gbps					
L0	5 (26) ¹	5 (26)	2.5 km	10 km	5 km	2 km	All	No
LE	11	16	10 km	N/A	10 km	5 km	3.x, 4.x	No

¹For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

Table 32 lists the extended ISL modes for switches that have a Condor ASIC.

Table 32 Extended ISL modes: switches with Condor ASIC.

Mode	Buffer Allocation			Distance at 1 Gbps	Distance at 2 Gbps	Distance at 4 Gbps	Earliest Fabric OS release	Extended fabrics license required?
	1 Gbps	2 Gbps	4 Gbps					
L0	5 (26) ²	5 (26)	2.5 km	10 km	5 km	2 km	All	No
LE	11	16	10 km	N/A	5 km	5 km	3.x, 4.x	No
L0.5	18	31	56	25 km	25 km	25 km	3.1.0, 4.1.0, 4.x, 5.x	Yes
L1	31	56	106	50 km	50 km	50 km	All	Yes
L12	56	106	206	100 km	100 km	100 km	All	Yes
LD ¹	Auto	Auto	Auto	Auto	Auto	Auto	3.1.0, 4.1.0, 4.4.0, 5.x (depending on the model)	Yes

¹The dynamic long-distance (LD) mode configures the number of buffer credits required, based on the actual link distance.

²For each data channel (in this case, there are 4) there are 5 credits, plus 6 extra credits.

For dynamic long-distance links, you can approximate the number of buffer credits using the following formula:

$$\text{Buffer credits} = [(\text{distance in km}) * (\text{data rate}) * 1000] / 2112$$

The data rate is 1.0625 for 1 Gbps, 2.125 for 2 Gbps, and 4.25 for 4 Gbps and Fibre Channel. This formula provides the minimum number of credits allowed to a given port; the actual number will likely be higher.

Configuring an extended ISL

Before configuring the ISL, ensure that the following conditions are met:

- Extended ISL support for HP StorageWorks 1 GB switches is limited as follows:
 - Extended ISLs are not supported between HP StorageWorks 1 GB switches and other HP StorageWorks models.
 - To support extended ISLs between HP StorageWorks 1 GB switches, the `fabric.ops.mode.longDistance` parameter must be set to 1 on all switches in the fabric. Each switch must be disabled before setting this parameter.
 - For fabrics that contain a mix of HP StorageWorks models, the `fabric.ops.mode.longDistance` parameter must be set to 0 (the default). Under certain circumstances (for example, if you want extended distance between Bloom-based switches) this mode needs to be enabled (set to 1) on switches running Fabric OS 3.x or 4.x. Talk to your switch provider for details.
 - The ports on both ends of the ISL must have the same configuration.
 - Use only qualified SFPs.
1. Connect to the switch and log in as admin.
 2. If the fabric contains a mix of switches, use the `configure` command to make sure the fabric-wide configuration parameter `fabric.ops.mode.longDistance` is set to 0 on all switches in the fabric. If the fabric contains HP StorageWorks 1-GB switches with extended ISLs, use the `switchDisable` command to disable the switch, and then use the `configure` command to set the fabric-wide configuration parameter `fabric.ops.mode.longDistance` to 1 on all switches in the fabric.

3. Issue the `portCfgLongDistance` command, using the following syntax:

```
portcfglongdistance [slotnumber/]portnumber [distance_level]
[vc_translation_link_init] [desired_distance]
```

where:

<i>slotnumber</i>	Specifies the slot number for Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director. This option is not applicable to fixed-port switches. The slot number must be followed by a slash (/) and the port number.
<i>portnumber</i>	Specifies the port number.
<i>distance_level</i>	Specifies the ISL mode to be set on the port (see Table 30 on page 164).
<i>vc_translation_link_init</i>	<p>Is an extended link initialization sequence, which is an enhanced link reset protocol, and prevents excessive resetting of ports.</p> <p>By default this option is set to 1 (enabled).</p> <p>To prevent fabric segmentation, this option must be set to 0 (disabled) when connecting to switches running Fabric OS versions earlier than 3.0.2c or 4.0.2.</p> <p>It must be set to 1 (enabled) when configuring a trunk over extended fabrics.</p>
<i>desired_distance</i>	<p>Is required for a port configured for LD mode. Specify the desired distance, in kilometers, for the link. The specified value is the upper limit for calculating buffer availability for the port. If the measured distance is more than the specified <i>desired_distance</i>, the port is allocated the number of buffers required by the specified desired distance. (Fabric OS versions earlier than 4.4.0 do not support this parameter.)</p> <p>Note that for a 4/256 SAN Director with FC4-16 or FC4-32 port blades, ports are shown as buffer-limited.</p>

4. Repeat [step 3](#) for the remote extended ISL port.

Both the local and remote extended ISL ports must be configured to the same distance level. When the connection is initiated, the fabric is reconfigured.

The following example configures slot 1, port 1 for the LD link distance mode, enables the extended link initialization sequence, and sets the desired distance to 50 kilometers:

```
switch:admin> portcfglongdistance 1/1 LD 1 50
switch:admin>
```

Trunking over distance

See ["Trunking over extended fabrics"](#) on page 173.

12 Administering ISL trunking

This chapter contains procedures for using the HP ISL Trunking licensed feature, which optimizes the use of bandwidth by allowing a group of ISLs to merge into a single logical link.

About ISL trunking

HP ISL Trunking reduces or eliminates situations that require static traffic routes and individual ISL management to achieve optimal performance. Trunking optimizes fabric performance by distributing traffic across the shared bandwidth of all the ISLs in a trunking group, allowing traffic to flow through any available link in a group rather than restricting it to a specific, potentially congested link. The use of trunking results in simplified fabric design and management, lowered cost of ownership, and increased data availability.

To use trunking, you must first install the HP ISL Trunking license. For details on obtaining and installing licensed features, see ["Maintaining licensed features"](#) on page 26. Trunking is enabled when the ISL Trunking license is activated and ports are reinitialized (after installing the license, you issue the `switchDisable` and `switchEnable` commands). Trunks are easily managed using either Fabric OS CLI commands or Advanced Web Tools. You can enable and disable trunking and set trunk port speeds (for example, 2 Gbit/sec, 4 Gbit/sec, or auto-negotiate) for entire switches or for individual ports.

Trunks distribute traffic dynamically and in order at the frame level, achieving greater performance with fewer ISLs.

Trunks are compatible with both short wavelength (SWL) and long wavelength (LWL) fiber optic cables and transceivers.

Figure 4 illustrates how trunking can result in more throughput by distributing data over four ISLs with no congestion. In a fabric that does not have trunking capability, some paths would be congested and other paths underutilized.

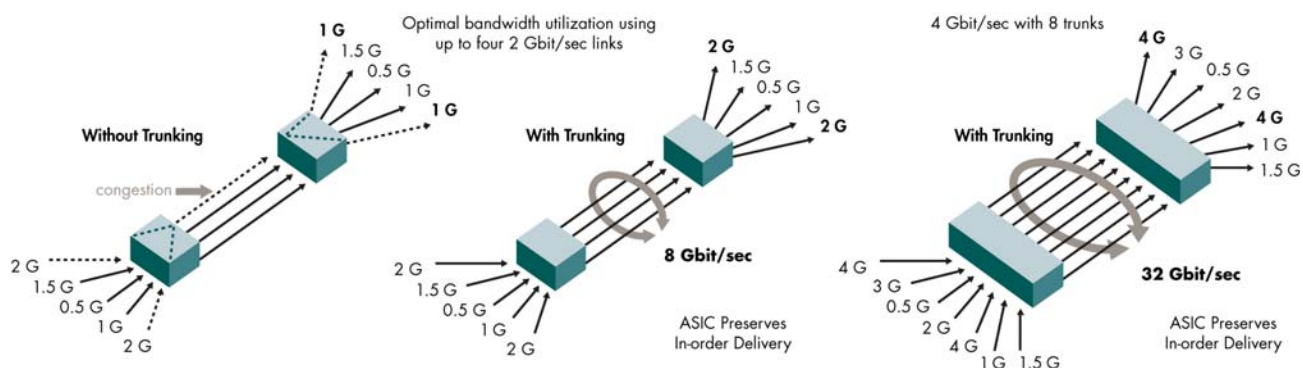


Figure 4 Distribution of traffic over ISL trunking groups

Trunks operate best when the cable length of each trunked link is roughly equal to the others in the trunk. Cable lengths for participating links should differ by no more than 550 meters. For optimal performance, HP recommends no more than 30 meters difference.

Connections between the SAN Switch 4/32 and 4/256 SAN Director (using FC4-16 and FC4-32 port blades) support these advanced features:

- Up to eight ports in one trunk group to create high-performance 32-Gbit/sec ISL trunks between switches
- ISL trunking over longer distances than other models
- Dynamic trunk master reassignment if a trunk master is disabled (on other platforms, all ports on a trunk must be disabled temporarily to reassign a master)
- 4 Gbit/sec trunk links

The maximum number of ports per trunk and trunks per switch depends on the HP StorageWorks model. For detailed information about trunking commands, see online help or the *HP StorageWorks Fabric OS 5.x command reference guide*.

Standard trunking criteria

Observe the following criteria for standard distance trunking:

- There must be a direct connection between participating switches.
- Trunk ports must reside in the same port group.
- Trunk ports must run at the same speed (either 2 Gbit/sec or 4 Gbit/sec).
- Trunk ports must be set to the same ISL mode (LO is the default). For details on extended ISL modes, see [Table 30](#) on page 164.
- Trunk ports must be E_Ports.
- Cable lengths for participating links should differ by no more than 550 meters.
- The `switch.interopMode` parameter must be set to 0. See "[Configuring interoperability mode](#)" on page 229 for information and procedures related to interoperability mode.
- The port ISL mode must be disabled (using the `portCfgIslMode` command).

Fabric considerations

The ISL trunking feature is provided with the Fabric OS and can be activated by entering a license key, available from the switch supplier. When the ISL Trunking license is activated (after you have entered the `switchDisable` and `switchEnable` commands), trunking is implemented for any eligible ISLs.

A license must be activated on each switch that participates in trunking. For the Core Switch 2/64, a single license key enables the feature on both logical switches.

To use ISL trunking in the fabric, the fabric must be designed to allow trunking groups to form. To identify the most useful trunking groups, evaluate the traffic patterns before designing or redesigning the fabric. This also applies to the SAN Director 2/128 configured with two domains, and the 4/256 SAN Director, which does not support two domains.

ISL Trunking can be used to simplify SAN design and improve performance. When designing the SAN, consider the following recommendations in addition to the standard guidelines for SAN design:

- Evaluate the traffic patterns within the fabric.
- Place trunking-capable switches adjacent to each other.
This maximizes the number of trunking groups that can form. If you are using a core/edge topology, place trunking-capable switches at the core of the fabric and place any switches that are not trunking-capable at the edge of the fabric.
- Activate an ISL Trunking license on each switch that is to participate in a trunking group.
- Cable lengths for participating links should differ by no more than 550 meters.
- When connecting two switches with two or more ISLs, ensure that all trunking requirements are met to allow a trunking group to form.
- Determine the optimal number of trunking groups between each set of linked switches, depending on traffic patterns and port availability.

The goal is to avoid traffic congestion without unnecessarily using ports that could be used to attach other switches or devices. Consider these points:

- Each physical ISL uses two ports that could otherwise be used to attach node devices or other switches.
- Trunking groups can be used to resolve ISL oversubscription if the total capability of the trunking group is not exceeded.

- Consider how the addition of a new path affects existing traffic patterns:
 - A trunking group has the same link cost as the master ISL of the group, regardless of the number of ISLs in the group. This allows slave ISLs to be added or removed without causing data to be rerouted, because the link cost remains constant.
 - The addition of a path that is shorter than existing paths causes traffic to be rerouted through that path.
 - The addition of a path that is longer than existing paths might not be useful because the traffic chooses the shorter paths first.
- Plan for future bandwidth addition to accommodate increased traffic.
For trunking groups over which traffic is likely to increase as business requirements grow, consider leaving one or two ports in the group available for future non-destructive addition of bandwidth.
- Consider creating redundant trunking groups where additional ports are available or paths are particularly critical.
This helps to protect against oversubscription of trunking groups, multiple ISL failures in the same group, and the rare occurrence of an ASIC failure.
- To provide the highest level of reliability, deploy trunking groups in redundant fabrics to further ensure ISL failures do not disrupt business operations.

Initializing trunking on ports

After you unlock the ISL Trunking license, you must reinitialize the ports being used for ISLs so that they recognize that trunking is enabled. This procedure needs to be performed only once.

To reinitialize the ports, you can either disable and then reenabling the switch, or disable and then reenabling the affected ports.

Disabling and reenabling the switch

1. Connect to the switch and log in as admin.
2. Issue the `switchDisable` command.
3. Issue the `switchEnable` command.

Disabling and reenabling ports

1. Connect to the switch and log in as admin.
2. Issue the `portDisable` command.

The format is:

```
portDisable [slot/]port
```

where *slot* is the slot number (Core Switch 2/64 and SAN Director 2/128 only) and *port* is the port number of the port you want to disable.

3. Issue the `portEnable` command.

The format is:

```
portEnable [slot/]port
```

where *slot* is the slot number (Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director only) and *port* is the port number of the port you want to enable.

Monitoring traffic

To implement ISL trunking effectively, you must monitor fabric traffic to identify congested paths or to identify frequently dropped links. While monitoring changes in traffic patterns, you can adjust the fabric design accordingly, such as by adding, removing, or reconfiguring ISLs and trunking groups in problem areas.

There are three methods of monitoring fabric traffic:

- Advanced Performance Monitoring monitors traffic flow and allows you to view the impact of different fabric configurations on performance. See ["Administering advanced performance monitoring"](#) on page 199.

- Fabric Watch allows you to monitor traffic flow through specified ports on the switch and send alerts when the traffic exceeds or drops below configurable thresholds. See the *HP StorageWorks Fabric OS 5.x Fabric Watch administrator guide* for additional information.
- Issue the `portPerfShow` command, as described in the following procedure, to record traffic volume for each port in your fabric over time.

Using the `portperfshow` command

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
portperfshow [interval]
```

where *interval* is the number of seconds between each data-gathering sample (the default is one sample every second).

3. Record the traffic flow for each port participating in an ISL.
4. Repeat [step 1](#) through [step 3](#) for each switch in the fabric until all ISL traffic flow is captured.

In a large fabric, it might be necessary to identify and capture only the key ISLs. However, you might want to continue this process throughout the day (or an entire work cycle), to capture varying traffic patterns under different conditions.

The following example shows a switch without trunking, and indicates that ports 0 through 2 are under utilized and ports 4 and 5 are congested:

```
switch:admin> portperfshow
0          1          2          3  4567 Total
-----
0          0          0          145m204m202m0168m 719
0          0          0          145m206m208m0186m 745
switch:admin>
```

The following example shows traffic flowing through a trunking group (ports 5, 6, and 7). After port 6 fails, traffic is redistributed over the remaining two links in the group, ports 5 and 7:

```
switch:admin> portperfshow
0          1          2          3  4567 Total
-----
0          0          0          0  0145m144m145m 434
0          0          0          0  0144m143m144m 431
0          0          0          0  0162m0162m 324
0          0          0          0  0186m0186m 372
0          0          0          0  0193m0192m 385
0          0          0          0  0202m0202m 404
0          0          0          0  0209m0209m 418
switch:admin>
```

Enabling and disabling ISL trunking

You can enable or disable HP ISL Trunking for a single port or for an entire switch. When you execute the `portCfgTrunkPort` or `switchCfgTrunk` command to update the trunking configuration, the ports for which the configuration applies are disabled and reenabled with the new trunk configuration. As a result, traffic through those ports could be disrupted.

Enabling or disabling ISL trunking on one port

1. Connect to the switch and log in as admin.
2. Issue the `portCfgTrunkPort` command.

The format is:

```
portcfgtrunkport [slotnumber/]portnumber mode
```

where:

<i>slotnumber</i>	Specifies the number of the slot in which the port blade containing the port is located (this operand is required only for switches with slots such as the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director).
<i>portnumber</i>	Specifies the number of the port on which you want to enable or disable trunking.
<i>mode</i>	Enables (1) or disables (0) trunking on the specified port.

The following example enables trunking on slot 1, port 3:

```
switch:admin> portcfgtrunkport 1/3 1
done.
switch:admin>
```

Enabling or disabling ISL trunking for all of the ports on a switch

1. Connect to the switch and log in as admin.
2. Issue the `switchCfgTrunk` command. The format is:

```
switchcfgtrunk mode
```

Mode 1 enables and mode 0 disables ISL Trunking for all ports on the switch.

The following example enables trunking all ports in the switch.

```
switch:admin> switchcfgtrunk 1
Committing configuration...done.
switch:admin>
```

Setting port speeds

For LD ports, if a port is set to auto-negotiate port speed, the maximum speed (which is 4 Gbit/sec) is assumed for reserving buffers for the port. This wastes buffers if the port is actually running at 2 Gbit/sec. For LD ports, it is best to set the port speed (this applies to the SAN Switch 4/32 and 4/256 SAN Director only).

You can set the port speed for one port or for an entire switch. Trunked ports must be set to the same speed.

Setting the speed for one port

1. Connect to the switch and log in as admin.

2. Issue the `portCfgSpeed` command:

```
portcfgspeed [slotnumber/]portnumber, speed_level
```

where:

<i>slotnumber</i>	Is for bladed systems only; it specifies the slot number of the port to be configured, followed by a slash (/). This operand is required only for switches with slots, such as the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director.
<i>portnumber</i>	Specifies the port number relative to its slot for bladed systems.
<i>speed_level</i>	Specifies the speed of the link: <ul style="list-style-type: none">• 0—Autonegotiating mode. The port configures for the highest speed.• 1—one Gbit/sec mode. Fixes the port at a speed of one Gbit/sec. Changing the speed to one Gbit/sec causes the port to be excluded from the trunk group.• 2—two Gbit/sec mode. Fixes the port at a speed of two Gbit/sec.• 4—four Gbit/sec mode. Fixes the port at a speed of four Gbit/sec. (4/8 SAN Switch, 4/16 SAN Switch, and SAN Switch 4/32, and 4/256 SAN Director only.)

The following example sets the speed for port 3 on slot 2 to two Gbit/sec:

```
switch:admin> portcfgspeed 2/3 2
done.
switch:admin>
```

The following example sets the speed for port 3 on slot 2 to auto-negotiate:

```
switch:admin> portcfgspeed s/3 0
done.
switch:admin>
```

Setting the speed for all of the ports on the switch

1. Connect to the switch and log in as admin.

2. Issue the `switchCfgSpeed` command. The format is:

```
switchcfgspeed speedlevel
```

where:

<i>speedlevel</i>	Specifies the speed of the link: <ul style="list-style-type: none">• 0—Auto-negotiating mode. The port configures for the highest speed.• 1—Fixes the port at a speed of 1 Gbit/sec. Changing the speed to 1 Gbit/sec causes the port to be excluded from the trunk group.• 2—Fixes the port at a speed of 2 Gbit/second.• 4—Fixes the port at a speed of 4 Gbit/second. (4/8 SAN Switch, 4/16 SAN Switch, and SAN Switch 4/32, and 4/256 SAN Director only.)
-------------------	--

The following example sets the speed for all ports on the switch to 2 Gbit/second:

```
switch:admin> switchcfgspeed 2
Committing configuration...done.
switch:admin>
```

The following example sets the speed for all ports on the switch to auto-negotiate:

```
switch:admin> switchcfgspeed 0
Committing configuration...done.
switch:admin>
```

Displaying trunking information

The `trunkShow` command offers an efficient means of listing all the trunks and members of a trunk. You can easily discover the peer ports for disabling a port, disconnecting a port, or adding additional members. Viewing trunks is also useful for when you want to ensure that trunks are formed correctly.

Use the `trunkShow` command to display the following information about ISL trunking groups:

- Number identifier.
 - Port-to-port connections, listed in the following format: *local_port_number -> remote_port_number*.
 - WWNs of the remote switches.
 - Deskew values (the time difference, in nanoseconds divided by 10, for traffic to travel over each ISL as compared to the shortest ISL in the group). The system sets the minimum deskew value of the shortest ISL to 15.
 - Master ports.
1. Connect to the switch and log in as admin.
 2. Issue the `trunkShow` command.

The following example shows three trunking groups (1, 2, and 3); ports 1, 4, and 14 are masters:

```
switch:admin> trunkshow
1: 1 -> 1    10:00:00:60:69:04:10:83    deskew 16 Master
    0 -> 0    10:00:00:60:69:04:10:83    deskew 15
2: 4 -> 4    10:00:00:60:69:04:01:94    deskew 16 Master
    5 -> 5    10:00:00:60:69:04:01:94    deskew 15
    7 -> 7    10:00:00:60:69:04:01:94    deskew 17
    6 -> 6    10:00:00:60:69:04:01:94    deskew 16
3:14 -> 14   10:00:00:60:69:04:10:83    deskew 16 Master
    15 -> 15   10:00:00:60:69:04:10:83    deskew 15
switch:admin>
```

Trunking over extended fabrics

In addition to the criteria listed in ["Standard trunking criteria"](#) on page 168, observe the following criteria for trunking over extended fabrics:

- ISL trunking over extended fabrics is supported on switches running Fabric OS 3.2.0 (or later) or 4.4.0 (or later).
- Extended Fabrics and ISL Trunking licenses are required on all participating switches.
- The `vc_translation_link_init` parameter must be set the same on all ports in an extended trunk. (For details on this parameter, see ["For dynamic long-distance links, you can approximate the number of buffer credits using the following formula:"](#) on page 165.)

Troubleshooting trunking problems

If you have difficulty with trunking, try the solutions in this section.

Listing link characteristics

If a link that is part of an ISL trunk fails, use the `trunkDebug` command to troubleshoot the problem, as shown in the following procedure:

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
trunkDebug port, port
```

where *port* specifies the number of a port in an ISL trunking group.

The `trunkDebug` command displays the possible reason that two ports cannot be trunked, including the following reasons:

- The switch does not support trunking.
- A trunking license is required.
- Trunking is not supported in switch interoperability mode.
- Port trunking is disabled.
- The port is not an E_Port.
- The port is not 2 Gbit/sec or 4 Gbit/sec.
- The port connects to different switches.
- The ports are not same speed, or they are not set to a valid speed.
- The ports are not set to the same LD mode.
- Local or remote ports are not in same port group.
- The difference in the cable length among trunked links is greater than the allowed difference.

The following example shows that port 3 is not configured as an E_Port:

```
switch:admin> trunkdebug 3 5
port 3 is not E port
switch:admin>
```

Recognizing buffer underallocation

For the SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128, if there is an underallocation or overcommitment of buffers to ports configured for extended trunking, the switches at both ends of the trunk try to disable some ports, so others can operate using the available buffers. (Standard trunks are not affected by buffer allocation.)

This issue of buffer underallocation does not apply to the SAN Switch 4/32 and 4/256 SAN Director.

A port disabled at one end because of buffer underallocation causes all the disabled ports at the other end to become enabled. Some of these enabled ports become disabled due to a lack of buffers, which in turn triggers ports to be enabled once again at the other end. While the system is stabilizing the buffer allocation, it warns that ports are disabled due to lack of buffers, but it does not send a message to the console when buffers are enabled. The system requires a few passes to stabilize the buffer allocation. Ultimately, the number of ports for which buffers are available come up and stabilize. Wait for stabilization, and then correct the buffer allocation situation.

Getting out of buffer-limited mode on E_Ports or LD_Ports:

1. Change the LD/L1/L2/L0.5 port speed to a lower speed (of non-buffer-limited ports).
2. Change the LD port's estimated distance to a shorter distance (of non-buffer-limited ports).
3. Change LD/L1/L2/L0.5 back to L0 (of non-buffer-limited ports).

4. If you are in buffer-limited mode on the LD port, increase the estimated distance.

These changes are implemented only after disabling (`portDisable`) and enabling (`portEnable`) the buffer-limited port (or buffer-limited switch).

Reconfiguring a port to LD from another mode can result in the port being disabled for lack of buffers—this does not apply to the SAN Switch 4/32 and 4/256 SAN Director (using FC4-16 and FC4-32 port blades). If this happens:

- In Fabric OS 4.2.x, reconfigure the disabled LD port back to the original mode.
- In Fabric OS 4.4.0 and later, specify a slightly shorter distance for the `desired_distance` parameter in the `portCfgLongDistance` command.

13 Administering advanced zoning

This chapter provides procedures for using the HP Advanced Zoning feature.

Zoning terminology

The following terms are used in the advanced zoning procedures:

- A *zone* is a region within the fabric where a specified group of fabric-connected devices (called *zone members*) have access to one another. When zoning is enabled, objects not explicitly defined in a zone are isolated, and members in the zoned fabric do not have access to them.
- A group of one or more zones is called a *zone configuration*.
- The complete set of all zone members defined in a fabric is called the *defined zone configuration*.
- Zoning procedures change zone objects in the defined configuration. When you enable a configuration with the `cfgEnable` command, it becomes the *effective zone configuration*. The effective zone configuration is restored after a switch reboot.
- A copy of the defined zone configuration (plus the name of the effective zone configuration) can be saved with the `cfgSave` command. The resulting *saved zone configuration* is restored after a switch reboot. If you make changes to the defined zone configuration but do not save them, the defined zone configuration and the saved zone configuration will be different.

Advanced zoning licenses are installed as active on all HP StorageWorks switches when they ship from the factory. They need to be reinstalled only if the license has been removed.

If a Zoning license is removed, you must make sure it is reinstalled properly on the affected switch before attempting the `cfgEnable` zoning operation. Failure to follow these steps can cause inconsistency of the zoning configuration on the affected switches if a zoning operation is attempted from a remote switch in the fabric. On the affected switches, an error message indicates that the Zoning license is missing.

You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions, for example, to create a temporary zone to back up non-member devices.

Any zone object connected to the fabric can be included in one or more zones. Zone objects can communicate only with other objects in the same zone. For example, consider [Figure 5](#), which shows:

- Three zones are configured, named Red, Green, and Blue.
- Server 1 can communicate only with the Loop 1 devices.
- Server 2 can communicate only with the RAID and Blue zone devices.
- Server 3 can communicate with the RAID device and the Loop1 device.
- The Loop 2 JBODs are not assigned to a zone; no other zoned fabric device can access them.

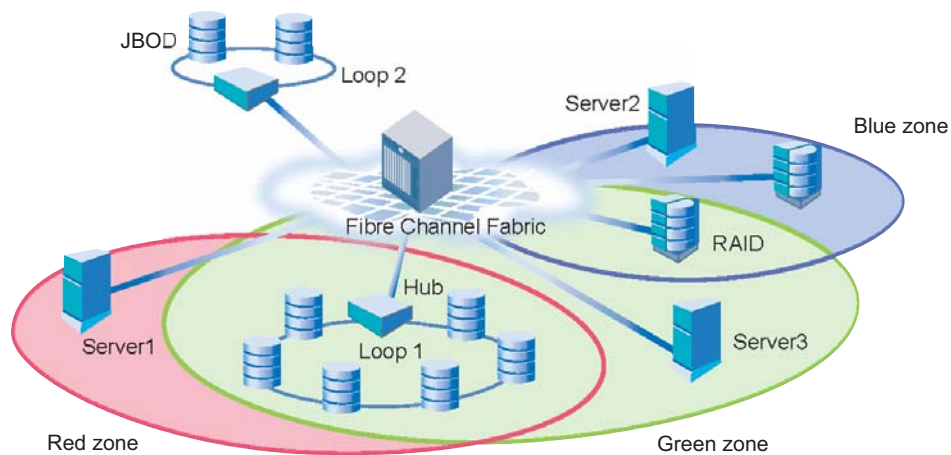


Figure 5 Zoning example

To list the commands associated with zoning, use the `zoneHelp` command. For detailed information on the zoning commands used in the procedures, see the *HP StorageWorks Fabric OS 5.x command reference guide* or to the online man page for each command.

Zoning concepts

Before using the procedures, become familiar with the zoning concepts described in the following sections.

Zone types

[Table 33](#) summarizes the types of zoning.

Table 33 Types of zoning

Zone type	Description
Storage-based	Storage units typically implement logical unit number (LUN)-based zoning, also called <i>LUN masking</i> . LUN-based zoning limits access to the LUNs on the storage port to the specific WWN of the server HBA. It is needed in most SANs. It functions during the probe portion of the SCSI initialization. The server probes the storage port for a list of available LUNs and their properties. The storage system compares the WWN of the requesting HBA to the defined zone list, and returns the LUNs assigned to the WWN. Other LUNs on the storage port are not made available to the server.
Host-based	Host-based zoning can implement WWN or LUN masking.
Fabric-based	<p>Fabric switches implement fabric-based zoning, in which the zone members are identified by WWN or port location in the fabric. Fabric-based zoning is also called <i>name server-based</i> or <i>soft</i> zoning.</p> <p>HP StorageWorks switches might also provide additional hardware enforcement of the zone. When a device queries the fabric Name Server, the Name Server determines the zones in which the device belongs. The server returns information on all members of the zones in the fabric to the device. Devices in the zone are identified by node WWN, port WWN, or domain, port of the switch to which the device is connected.</p> <p>Fabric-based zoning is perhaps the most controversial aspect of zoning. There are several approaches for implementing fabric zoning; all of them work in most cases. However, there are pros and cons to each form. The primary forms are summarized in Table 34.</p>

Table 34 Approaches to fabric-based zoning

Form	Description
Single HBA	Zoning by single HBA most closely re-creates the original SCSI bus. Each zone created has only one HBA (initiator) in the zone; each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports are added. If the HBA also accesses tape devices, a second zone is created with the HBA and associated tape devices in it. In the case of clustered systems, it could be appropriate to have an HBA from each of the cluster members included in the zone; this is equivalent to having a shared SCSI bus between the cluster members and presumes that the clustering software can manage access to the shared devices. In a large fabric, zoning by single HBA requires the creation of possibly hundreds of zones; however, each zone contains only a few members. Zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change. This zoning philosophy is the preferred method.
Application	Zoning by application typically requires zoning multiple, perhaps incompatible, operating systems into the same zones. This method of zoning creates the possibility that a minor server in the application suite could disrupt a major server (such as a web server) disrupting a data warehouse server. Zoning by application can also result in a zone with a large number of members, providing greater susceptibility to administrative errors, such as RSCNs going out to a larger group than necessary.
Operating system	Zoning by operating system has issues similar to zoning by application. In a large site, this type of zone can become very large and complex. When zone changes are made, they typically involve applications rather than a particular server type. If members of different operating system clusters can see storage assigned to another cluster, they might attempt to own the other cluster's storage and compromise the stability of the clusters.
Port allocation	Avoid zoning by port allocation unless the administration team has very rigidly enforced processes for port and device allocation in the fabric. It does, however, provide some positive features. For instance, when a storage port, server HBA, or tape drive is replaced, the change of WWN for the new device is of no consequence. As long as the new device is connected to the original port, it continues to have the same access rights. The ports on the edge switches can be pre-associated to storage ports, and control of the fan-in ratio (the ratio of the input port to output port) can be established. With this pre-assigning technique, the administrative team cannot overload any one storage port by associating too many servers with it.
No fabric zoning	Using no fabric zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should be used only in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

Zone objects

A *zone object* is any device in a zone, such as the:

- Physical port number or area ID on the switch
- Node WWN (N-WWN)
- Port WWN (P-WWN)

Zone objects identified by port number or area number are specified as a pair of decimal numbers in the form *d, area* (*d* is the domain ID of the switch and *area* is the area number on that switch).

For example, on Core Switch 2/64 or SAN Director 2/128, 4, 46 specifies port 14 in slot number 3 (domain ID 4, area 46). On fixed-port models, 3, 13 specifies port 13 in switch domain ID 3.

When the physical port number specifies a zone object, all devices connected to that port are in the zone. If the physical port is an arbitrated loop, all devices on the loop are part of the zone.

WWNs are specified as 8-byte (16-digit) hexadecimal numbers, separated by colons, for example, 10:00:00:90:69:00:00:8a. When a node name specifies a zone object, all ports on such a device are in the zone. When a port name specifies a zone object, only the single port is in the zone.

The types of zone objects used to define a zone can be mixed and matched. For example, a zone defined with the zone objects 2,12; 2,14; 10:00:00:80:33:3f:aa:11 contains the devices connected to domain 2, ports 12 and 14, and a device with the WWN (either node name or port name) 10:00:00:80:33:3f:aa:11 that is connected on the fabric.

Zone aliases

A *zone alias* is a name assigned to a device or a group of devices. By creating an alias, you can assign a familiar name to a device or group multiple devices into a single name. This simplifies cumbersome data entry and allows an intuitive naming structure (such as using NT_Hosts to define all NT hosts in the fabric).

Zone aliases also simplify repetitive entry of zone objects, such as port numbers or a WWN. For example, you can use the name Eng as an alias for 10:00:00:80:33:3f:aa:11.

A useful convention is to name zones for the initiator they contain. For example, if you use the alias SRV_MAILSERVER_SLT5 to designate a mail server in PCI slot 5, the alias for the associated zone is ZNE_MAILSERVER_SLT5. This clearly identifies the server HBA associated with the zone.

Zone configuration naming is more flexible. One configuration should be named *PROD_fabricname*, where *fabricname* is the name that the fabric has been designated. The purpose of the PROD configuration is to easily identify the configuration that can be implemented and provide the most generic services. If other configurations are used for specialized purposes, names such as *BACKUP_A*, *RECOVERY_2*, and *TEST_18jun02* can be used.

Zone configurations

A *zone configuration* is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is in effect, all zones that are members of that configuration are in effect.

The different types of zone configurations are:

- **Defined configuration:** The complete set of all zone objects defined in the fabric.
- **Effective configuration:** A single zone configuration that is currently in effect. The effective configuration is built when an administrator enables a specified zone configuration.
- **Saved configuration:** A copy of the defined configuration plus the name of the effective configuration, which is saved in flash memory by the `cfgSave` command. (You can also use the `configUpload` command to provide a backup of the zoning configuration and the `configDownload` command to restore the zoning configuration.) There might be differences between the saved configuration and the defined configuration if the system administrator has modified any of the zone definitions and has not saved the configuration.
- **Disabled configuration:** The effective configuration is removed from flash memory.

On power-up, the switch reloads the saved configuration. If a configuration was active when it was saved, the same configuration is reinstated on the local switch with an `autorun` of the `cfgEnable` command.

You can establish a zone by identifying zone objects using one or more of the following zoning schemes:

- **Domain, port number level:** All members are specified by domain ID, port number, or domain, area number pair or aliases, described in “[Zone aliases](#)” on page 180.
- **WWN level:** All members are specified only by WWNs or aliases of WWNs. They can be node or port versions of the WWN.
- **Mixed zoning:** A zone containing members specified by a combination of *domain, port number*, and/or *domain, area number* and WWN.

Zoning enforcement

Software-enforced and hardware-enforced zoning are supported.

Software-enforced zoning

Zoning enables users to restrict access to devices in a fabric. Software-enforced zoning prevents hosts from discovering unauthorized target devices, while hardware-enforced zoning prevents a host from accessing a device it is not authorized to access.

Software-enforced zoning:

- Is also called *soft zoning*, *Name Server zoning*, *fabric-based zoning*, *session-based zoning*, or *hardware-assisted zoning*.
- Is available on 1-Gbit/sec, 2-Gbit/sec, and 4-Gbit/sec platforms.
- Prevents hosts from discovering unauthorized target devices.
- Ensures that the Name Server does not return any information to an unauthorized initiator in response to a Name Server query.
- Is always active whenever a zone configuration is in effect.
- Does not prohibit access to the device. If an initiator has knowledge of the network address of a target device, it does not need to query the Name Server to access it, which could lead to undesired access to a target device by unauthorized hosts.
- Is exclusively enforced through selective information presented to end nodes through the fabric SNS. When an initiator queries the Name Server for accessible devices in the fabric, the Name Server returns only those devices that are in the same zone as the initiator. Devices that are not part of the zone are not returned as accessible devices.

Hardware-enforced zoning

Hardware-enforced zoning is specified without using the mixed-zoning scheme (mixed zones contain domains, ports and WWNs as zone members). HP StorageWorks switches augment software-enforced zoning with hardware enforcement. The exact methodology varies on different switch models.

Hardware-enforced zoning (also called *hard zoning*):

- Prevents a host from accessing a device it is not authorized to access.
- Checks each frame before it is delivered to a zone member and discards it if there is a zone mismatch. When hardware-enforced zoning is active, the switch monitors the communications and blocks any frames that do not comply with the effective zone configuration. The switch performs this blocking at the transmit side of the port on which the destination device is located.
- Is enforced at the ASIC level. Each ASIC maintains a list of source port IDs that have permission to access any of the ports on that ASIC.

Fabric OS uses hardware-enforced zoning (on a per-zone basis) whenever the fabric membership or zone configuration changes.

Table 35 shows the various HP StorageWorks switch models, the hardware zoning methodology for each, and tips for best usage.

Table 35 Enforcing hardware zoning

Fabric type	Methodology	Best practice
HP StorageWorks 1-GB switches	<p>Enables hardware-enforced zoning only on domain, port zones; WWN or mixed zones are not hardware-enforced. Any domain, port zone that overlaps a mixed or WWN zone is not hardware-enforced.</p> <p>An overlap occurs when a member specified by WWN is connected to a port in a domain, port zone. The domain, port zone loses its hardware enforcement even though a review of the zone configuration does not indicate it.</p>	Use domain, port identifiers (PIDs). Do not identify a zone member by its WWN.
4/8 SAN Switch, 4/16 SAN Switch, HP StorageWorks 2-GB switches, SAN Switch 4/32, Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director	<p>Enable hardware-enforced zoning on domain, port zones, and WWN zones. Overlap of similar zone types does not result in the loss of hardware enforcement. Overlap with other zone type results in the loss of hardware enforcement.</p> <p>As in the HP StorageWorks 1 GB switches, connecting a device specified by WWN into a port specified in a domain, port zone results in loss of the hardware enforcement in both zones.</p>	Use either WWN or domain, port identifiers.
Mixed switches	<p>Enable hardware-enforced zoning according to each switch type. Use the <code>portZoneShow</code> command to find the zone type to which a device is attached.</p>	<p>Use domain, port identifiers.</p> <p>You can use WWN identifiers if you place disk and tape targets on HP StorageWorks 2-GB switches, Core Switch 2/64, and SAN Director 2/128, and do not use domain, port identifiers.</p>

Figure 6 shows a fabric with four non-overlap ping hardware-enforced zones.

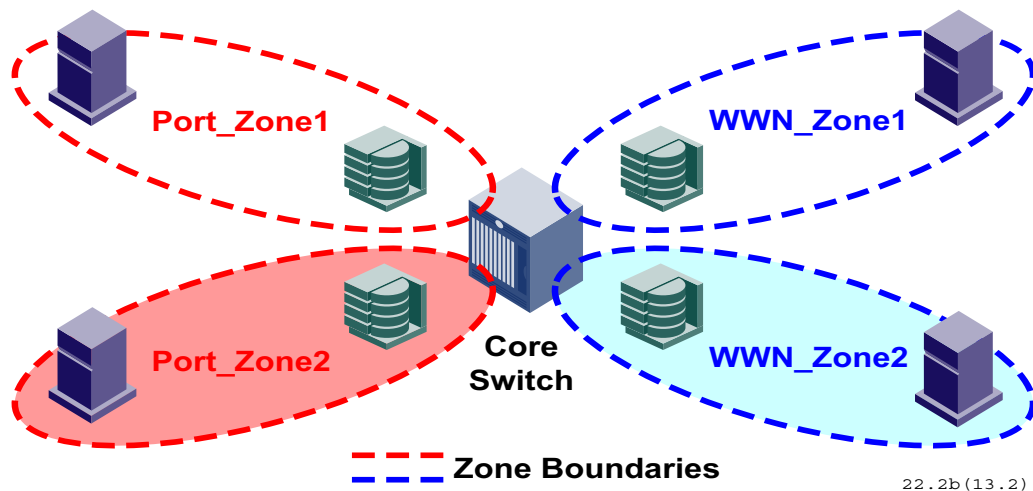


Figure 6 Hardware-enforced non-overlapping zones

Figure 7 shows the same fabric components zoned in an overlapping fashion.

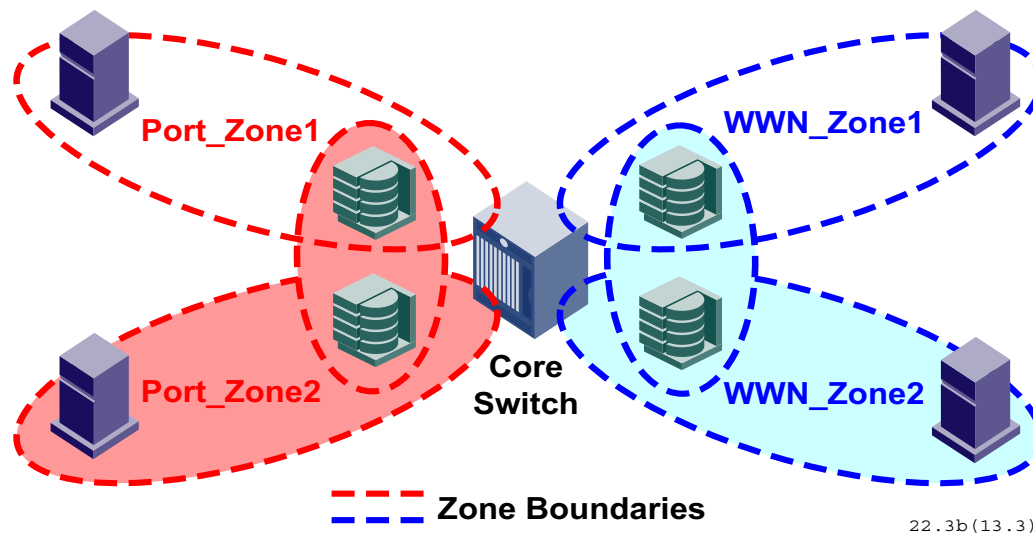
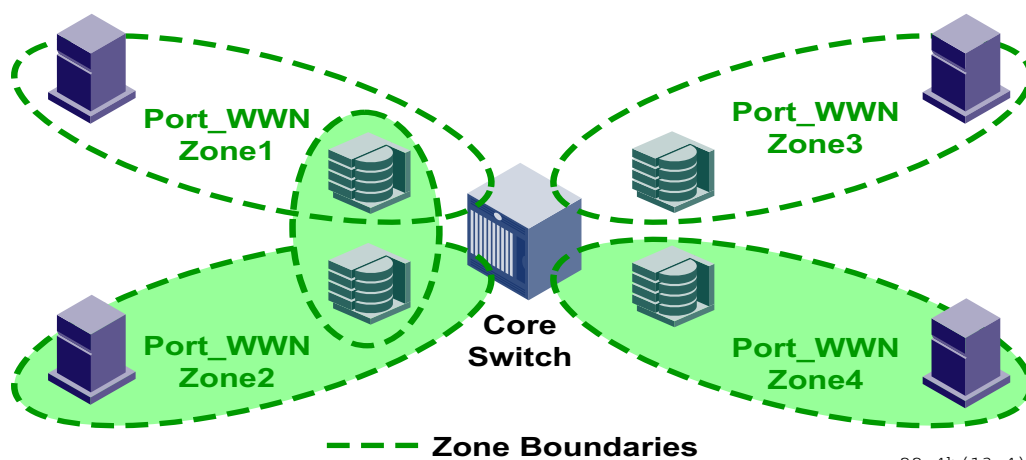


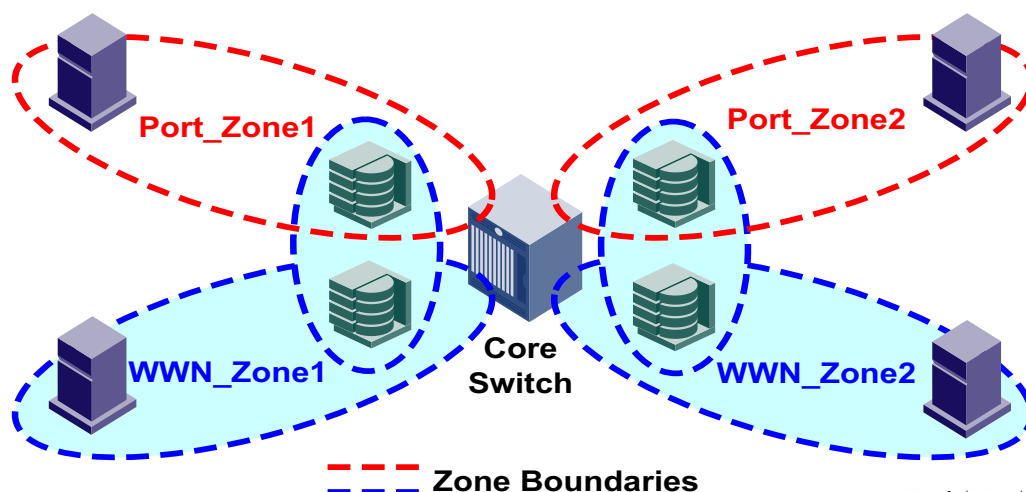
Figure 7 Hardware-enforced overlapping zones

Any zone using both WWNs and domain, port entries on the HP StorageWorks 2-Gbit/sec platform relies on Name Server authentication as well as hardware-assisted (ASIC) authentication, which ensures that any PLOGI/ADISC/PDISC/ACC from an unauthorized device attempting to access a device it is not zoned with is rejected. HP StorageWorks 2-Gbit/sec switches always deploy the hardware assist, in any zone configuration (see [Figure 8](#) and [Figure 9](#)).



22.4b (13.4)

Figure 8 Zoning with hardware assist (mixed-port and WWN zones)



22.5b (13.5)

Figure 9 Session-based hard zoning

In **Figure 9**, only the ports that are overlapped are software-enforced with hardware assist.


Rules for configuring zones

Observe the following rules when configuring zones.

- If security is a priority, use hard zoning.
- The use of aliases is optional with zoning, and using aliases requires structure when defining zones. However, aliases aid administrators of a zoned fabric to understand the structure and context.
- Evaluate the security requirements of the fabric. If additional security is required, add HP Secure Fabric OS into the fabric.
- If the fabric includes an HP StorageWorks switch and you support a third-party switch product, they are able to use only WWN zoning; other types of zoning, including QuickLoop, are not supported.

- **QuickLoop:** Evaluate whether the fabric will also use QuickLoop Fabric Assist (QLFA) or QuickLoop. If you are running Fabric OS 4.x, consider the following before creating and setting up QLFA zones:
 - **QuickLoop Zoning.** QuickLoop and QuickLoop zones cannot run on switches running Fabric OS 4.x. However, Fabric OS 4.x can still manage (create, remove, update) QuickLoop zones on any non-4.x switch.
 - **QLFA.** Fabric OS 4.x cannot have a Fabric Assist host directly connected to it. However, targets on a Fabric OS 4.x switch can still be part of a Fabric Assist zone if a Fabric Assist host is connected to a non-4.x switch.
- **Zone changes:** Zone changes in a production fabric can cause a disruption of I/O when an RSCN is generated because of the zone change and the HBA is unable to process the RSCN fast enough. Although RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, perform zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can rectify the situation.
- **Final verification:** After changing or enabling a zone configuration, confirm that the nodes and storage devices can identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.

The zone configuration is managed on a fabric basis. Zoning can be implemented and administered from any switch in the fabric that has an Advanced Zoning license enabled. When a change in the configuration is saved, enabled, or disabled per the transactional model, it is distributed (by closing the transaction) to all switches in the fabric, preventing a single point of failure for zone information.

 **NOTE:** Zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent commands. For a large fabric, you might want to wait several minutes between commands.

Creating and managing zone aliases

A zone alias is a logical group of ports, WWNs, or arbitrated loop physical addresses (AL_PAs). You can simplify the process of creating zones by first specifying aliases, which eliminates the need for long lists of individual zone-member names.

Note that if you are creating a new alias using `aliCreate w, "1,1"`, and a user in another telnet session executes `cfgEnable` (or `cfgDisable` or `cfgSave`), the other user's transaction aborts your transaction and you receive an error message. Creating a new alias while there is a zone merge taking place might also abort your transaction. For details about zone merging and zone merge conflicts, see ["Adding a new switch or fabric"](#) on page 195.

Creating an alias

1. Connect to the switch and log in as admin.
2. Issue the `aliCreate` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> alicreate "array1", "2,32; 2,33; 2,34; 4,4"
switch:admin> alicreate "array2", "21:00:00:20:37:0c:66:23; 4,3"
switch:admin> alicreate "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Adding members to an alias

1. Connect to the switch and log in as admin.
2. Issue the `aliAdd` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:


```
switch:admin> aliadd "array1", "1,2"
switch:admin> aliadd "array2", "21:00:00:20:37:0c:72:51"
switch:admin> aliadd "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Removing members from an alias

1. Connect to the switch and log in as admin.
2. Issue the `aliRemove` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> aliremove "array1", "1,2"
switch:admin> aliremove "array2", "21:00:00:20:37:0c:72:51"
switch:admin> aliremove "loop1", "4,6"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

 **NOTE:** For Fabric OS versions earlier than 4.4.0, when using the `aliRemove` command, the order in which the members appear in the list is critical. For more information on this command, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

Deleting an alias

1. Connect to the switch and log in as admin.
2. Issue the `aliDelete` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> alidelete "array1"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Viewing an alias in the defined configuration

1. Connect to the switch and log in as admin.
2. Issue the `aliShow` command.

The following example shows all zone aliases beginning with `arr`:

```
switch:admin> alishow "arr*"
alias: array1  21:00:00:20:37:0c:76:8c
alias: array2  21:00:00:20:37:0c:66:23
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

Creating and maintaining zones

Before executing `cfgDisable`, `cfgEnable`, or `cfgSave` commands, execute the `rscsDisabled` command to check whether your fabric has Reliable Commit Service (RCS) enabled (`rscsDisabled=0`). If RCS is disabled (`rscsDisabled=1`), check for older switches in the fabric. After the older switches are upgraded, RCS is enabled by default.

RCS is available on all switch versions 4.1 and later. RCS guarantees that either all or none of the switches receive the new zone configuration. HP recommends that you use RCS to secure a reliable propagation of the latest zone configuration. If you use non-RCS mode, you must log in to every switch to monitor the status of the zone configuration.

Creating a zone

1. Connect to the switch and log in as admin.
2. Issue the `zoneCreate` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> zonecreate "greenzone", "2,32; 2,33; 2,34; 4,4"
switch:admin> zonecreate "redzone", "21:00:00:20:37:0c:66:23; 4,3"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Adding devices (members) to a zone

1. Connect to the switch and log in as admin.
2. Issue the `zoneAdd` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> zoneadd "greenzone", "1,2"
switch:admin> zoneadd "redzone", "21:00:00:20:37:0c:72:51"
switch:admin> zoneadd "bluezone", "4,6; 21:00:00:20:37:0c:66:23"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Removing devices (members) from a zone

1. Connect to the switch and log in as admin.
2. Issue the `zoneRemove` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> zonremove "greenzone", "1,2"
switch:admin> zonremove "redzone", "21:00:00:20:37:0c:72:51"
switch:admin> zonremove "bluezone", "4,6; 21:00:00:20:37:0c:66:23
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Deleting a zone

1. Connect to the switch and log in as admin.
2. Issue the `zoneDelete` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> zonedel "bluezone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Viewing a zone in the defined configuration

1. Connect to the switch and log in as admin.
2. Issue the `zoneShow` command.

The following example shows all zones beginning with A, B, or C:

```
switch:admin> zoneshow "[A-C]*"
zone: Blue_zone 1,1; array1; 1,2; array2
zone: Bobs_zone 4,5; 4,6; 4,7; 4,8; 4,9
```

If no parameters are specified, the entire zone database (both the defined and effective configuration) is displayed.

Merging zones

Before linking two switches together, it is important that you know the zone database limit of adjacent switches. For example, when switches running Fabric OS 3.2, 4.4.0, or 5.x discover that the zone merge database is larger than its pre-determined zone database limit, they issue a reject notification before symmetrically segmenting their own ends of the ISL, thereby preventing the new switch from joining the fabric.

Symmetrical segmentation occurs when both ends of an ISL are shut down. Subsequently, no frames are exchanged between the two switches.

Asymmetrical segmentation not only prevents frames from being exchanged between switches, but also causes routing inconsistencies.

The best way to avoid either type of segmentation is to know the zone database size limit of adjacent switches. [Table 36](#) through [Table 39](#) provide the expected behavior based on different database sizes after a zone merge is specified.

Table 36 Resulting database size: 0 to 96K

Receiver → Initiator ↓	FOS 2.6	FOS 3.1	FOS 3.2	FOS 4.0, 4.1, 4.2	FOS 4.3, 4.4.0	FOS 5.5.0, 5.0.1	Fibre Channel Router	XPath 7.3
FOS 2.6, 3.1	Join	Join	Join	Join	Join	Join	Join	Join
FOS 3.2	Join	Join	Join	Join	Join	Join	Join	Join
FOS 4.0, 4.1, 4.1	Join	Join	Join	Join	Join	Join	Join	Join
FOS 4.3, 4.4.0	Join	Join	Join	Join	Join	Join	Join	Join
FOS 5.5.0, 5.0.1	Join	Join	Join	Join	Join	Join	Join	Join
Fibre Channel Router	Join	Join	Join	Join	Join	Join	Join	Join
XPath 7.3	Join	Join	Join	Join	Join	Join	Join	Join

Table 37 Resulting database size: 96K to 128K

Receiver → Initiator ↓	FOS 2.6	FOS 3.1	FOS 3.2	FOS 4.0, 4.1, 4.2	FOS 4.3, 4.4.0	FOS 5.5.0, 5.0.1	Fibre Channel Router	XPath 7.3
FOS 2.6, 3.1	Segment	Segment	Segment	Segment	Segment	Segment	Join	Segment
FOS 3.2	Segment	Segment	Join	Join	Join	Join	Join	Join
FOS 4.0, 4.1, 4.1	Segment	Segment	Segment	Join	Join	Join	Join	Join
FOS 4.3, 4.4.0	Segment	Segment	Join	Join	Join	Join	Join	Join
FOS 5.5.0, 5.0.1	Segment	Segment	Join	Join	Join	Join	Join	Join
Fibre Channel Router	Join	Join	Join	Join	Join	Join	Join	Join
XPath 7.3	Segment	Segment	Segment	Join	Join	Join	Join	Join

Table 38 Resulting database size: 128K to 256K

Receiver → Initiator ↓	FOS 2.6	FOS 3.1	FOS 3.2	FOS 4.0, 4.1, 4.2	FOS 4.3, 4.4.0	FOS 5.5.0, 5.0.1	Fibre Channel Router	XPath 7.3
FOS 2.6, 3.1	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS 3.2	Segment	Segment	Join	Segment	Join	Join	Join	Segment
FOS 4.0, 4.1, 4.1	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS 4.3, 4.4.0	Segment	Segment	Join	Segment	Join	Join	Join	Segment
FOS 5.5.0, 5.0.1	Segment	Segment	Join	Segment	Join	Join	Join	Segment
Fibre Channel Router	Join	Join	Join	Segment	Join	Join	Join	Segment
XPath 7.3	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment

Table 39 Resulting database size: 256K to 1M

Receiver → Initiator ↓	FOS 2.6	FOS 3.1	FOS 3.2	FOS 4.0, 4.1, 4.2	FOS 4.3, 4.4.0	FOS 5.5.0, 5.0.1	Fibre Channel Router	XPath 7.3
FOS 2.6, 3.1	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS 3.2	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS 4.0, 4.1, 4.1	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS 4.3, 4.4.0	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment
FOS 5.5.0, 5.0.1	Segment	Segment	Segment	Asymmetrical Segment	Segment	Join	Join	Segment
Fibre Channel Router	Segment	Segment	Segment	Segment	Segment	Join	Join	Segment
XPath 7.3	Segment	Segment	Segment	Segment	Segment	Segment	Segment	Segment

Creating and modifying zoning configurations

You can store a number of zones in a zoning configuration database. The maximum number of items that can be stored in the zoning configuration database depends on the following criteria:

- Number of switches in the fabric.
- Whether or not interoperability mode is enabled.
- Number of bytes per item. The number of bytes required for an item depends on the specifics of the fabric, but cannot exceed 64 bytes per item.

When enabling a new zone configuration, you must ensure that the size of the configuration does not exceed the minimum size supported by all switches in the fabric. This is particularly important if and when you downgrade to a Fabric Operating System (FOS) that supports a smaller zone database than the current FOS. In this scenario, the zone database in the current FOS would have to be changed to the smaller zone database before the downgrade.

You can use the `cfgSize` command to check both the maximum available size and the currently saved size on all switches. See the *HP StorageWorks Fabric OS 5.x command reference guide* for details on the `cfgSize` command. If you believe you are approaching the maximum, you can save a partially completed zoning configuration and use the `cfgSize` command to determine the remaining space.

For important considerations for managing zoning in a fabric, and details about the maximum zone database size for each version of the FOS, see "[Maintaining zone objects](#)" on page 193.

Creating a zoning configuration

1. Connect to the switch and log in as admin.
2. Issue the `cfgCreate` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> cfgcreate "NEW_cfg", "redzone; bluezone; greenzone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Adding zones (members) to a zoning configuration

1. Connect to the switch and log in as admin.
2. Issue the `cfgAdd` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> cfgadd "newcfg", "bluezone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Removing zones (members) from a zone configuration

1. Connect to the switch and log in as admin.
2. Issue the `cfgRemove` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> cfgremove "newcfg", "redzone"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Deleting a zone configuration

1. Connect to the switch and log in as admin.
2. Issue the `cfgDelete` command.
3. Issue the `cfgSave` command to save the change to the defined configuration.

For example:

```
switch:admin> cfgdelete "testcfg"
switch:admin> cfgsave
You are about to save the Defined zoning configuration. This
action will only save the changes on the Defined configuration.
Any changes made on the Effective configuration will not
take effect until it is re-enabled.
Do you want to save Defined zoning configuration only? (yes, y, no, n): [no] y
```

Clearing changes to a configuration

Use the `cfgTransAbort` command. When this command is executed, all changes since the last save operation (performed with the `cfgSave` command) are cleared.

In the following example, assume that the removal of a member from `zone1` was done in error:

```
switch:admin> zoneremove "zone1","3,5"
switch:admin> cfgtransabort
```

Viewing all zone configuration information

If you do not specify an operand when executing the `cfgShow` command to view zone configurations, all zone configuration information (both defined and effective) is displayed. If there is an outstanding transaction, the newly edited zone configuration that has not yet been saved, is displayed. If there are no outstanding transactions, the committed zone configuration is displayed.

1. Connect to the switch and log in as admin.
2. Issue the `cfgShow` command with no operands.

For example:

```
switch:admin> cfgshow
Defined configuration:
  cfg:   USA1      Blue_zone
  cfg:   USA_cfg Red_zone; Blue_zone
  zone:  Blue_zone
         1,1; array1; 1,2; array2
  zone:  Red_zone
         1,0; loop1
  alias: array1   21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
  alias: array2   21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
  alias: loop1    21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df

Effective configuration:
  cfg:   USA_cfg
  zone:  Blue_zone
         1,1
         21:00:00:20:37:0c:76:8c
         21:00:00:20:37:0c:71:02
         1,2
         21:00:00:20:37:0c:76:22
         21:00:00:20:37:0c:76:28
  zone:  Red_zone
         1,0
         21:00:00:20:37:0c:76:85
         21:00:00:20:37:0c:71:df
```

Viewing selected zone configuration information

1. Connect to the switch and log in as admin.
2. Issue the `cfgShow` command and specify a pattern.

For example, to display all zone configurations that start with `Test`:

```
switch:admin> cfgshow "Test*"
cfg:   Test1 Blue_zone
cfg:   Test_cfg Red_zone; Blue_zone
```

Viewing a configuration in the effective zone database

1. Connect to the switch and log in as admin.
2. Issue the `cfgActvShow` command.

For example:

```
switch:admin> cfgactvshow
Effective configuration:
  cfg:    NEW_cfg
  zone:   Blue_zone
    1,1
    21:00:00:20:37:0c:76:8c
    21:00:00:20:37:0c:71:02
    1,2
    21:00:00:20:37:0c:76:22
    21:00:00:20:37:0c:76:28
  zone:   Red_zone
    1,0
    21:00:00:20:37:0c:76:85
    21:00:00:20:37:0c:71:df
```

Maintaining zone objects

You can use the `cfgDelete` command to delete a zone configuration, but there is a quicker and easier way to perform the same task via the zone object commands (`zoneObjectExpunge`, `zoneObjectCopy`, and `zoneObjectRename`). You can also copy and rename zone objects. When you copy a zone object, the resulting object has the same type as the original. Deleting a zone object also removes the object from any member lists of other objects. You can rename objects for all zone object types.

Copying a zone object

1. Connect to the switch and log in as admin.
2. Issue the `cfgShow` command to view the zone configuration objects you want to copy.

For example, to display all zone configuration objects that start with `Test`:

```
switch:admin> cfgshow "Test*"
cfg:    Test1 Blue_zone
cfg:    Test_cfg Red_zone; Blue_zone
```

3. Issue the `zoneObjectCopy` command, specifying the zone configuration objects you want to copy, along with the new object name.

Note that zone configuration names are case-sensitive; blank spaces are ignored. For example:

```
switch:admin> zoneobjectcopy "Test1", "US_Test1"
```

4. Issue the `cfgShow` command to verify the new zone object is present.

For example:

```
switch:admin> cfgshow "Test*"
cfg:    Test1 Blue_zone
cfg:    Test_cfg Red_zone; Blue_zone
cfg:    US_Test1 Blue_zone
```

5. If you want the change preserved when the switch reboots, save it to nonvolatile memory (also known as *flash memory*) by issuing the `cfgSave` command.
6. For the change to become effective, enable the appropriate zone configuration using the `cfgEnable` command.

Deleting a zone object

1. Connect to the switch and log in as admin.
2. Issue the `cfgShow` command to view the zone configuration objects you want to delete.

For example:

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Red_zone; White_zone; Blue_zone
zone: Blue_zone
      1,1; array1; 1,2; array2
zone: Red_zone
      1,0; loop1
zone: White_zone
      1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Issue the `zoneObjectExpunge` command to delete the zone object.

Note that zone configuration names are case-sensitive; blank spaces are ignored. For example:

```
switch:admin> zoneObjectExpunge "White_zone"
```

4. Issue the `cfgShow` command to verify the deleted zone object is no longer present.
5. If you want the change preserved when the switch reboots, save it to nonvolatile memory (also known as *flash memory*) by issuing the `cfgSave` command.
6. For the change to become effective, enable the appropriate zone configuration using the `cfgEnable` command.

Renaming a zone object

1. Connect to the switch and log in as admin.
2. Issue the `cfgShow` command to view the zone configuration objects you want to rename.

For example:

```
switch:admin> cfgShow
Defined configuration:
cfg: USA_cfg Red_zone; White_zone; Blue_zone
zone: Blue_zone
      1,1; array1; 1,2; array2
zone: Red_zone
      1,0; loop1
zone: White_zone
      1,3; 1,4
alias: array1 21:00:00:20:37:0c:76:8c; 21:00:00:20:37:0c:71:02
alias: array2 21:00:00:20:37:0c:76:22; 21:00:00:20:37:0c:76:28
alias: loop1 21:00:00:20:37:0c:76:85; 21:00:00:20:37:0c:71:df
```

3. Issue the `zoneObjectRename` command to rename zone configuration objects.

Note that zone configuration names are case-sensitive; blank spaces are ignored. For example:

```
switch:admin> zoneObjectRename "White_zone", "Red_zone"
```

4. Issue the `cfgShow` command to verify the renamed zone object is present.
5. If you want the change preserved when the switch reboots, save it to nonvolatile (also known as *flash memory*) by issuing the `cfgSave` command.
6. For the change to become effective, enable the appropriate zone configuration using the `cfgEnable` command.

For details about the `zoneObjectCopy`, `cfgShow`, `cfgEnable`, and `cfgSave` commands, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

Managing zoning configurations in a fabric

To modify an existing zone configuration, you can add, delete, or remove individual elements to create the desired configuration. After the changes have been made, save the configuration to ensure the configuration is permanently saved in the switch and that the configuration is replicated throughout the fabric.

The switch configuration file can also be uploaded to the host for archiving; it can also be downloaded from the host to a switch in the fabric. See "[Backing up a configuration](#)" on page 73, "[Restoring a configuration](#)" on page 74, or the `configUpload` and `configDownload` commands in the *HP StorageWorks Fabric OS 5.x command reference guide*.

[Table 40](#) presents zoning database size limitations for various Fabric OS release versions. The maximum size of a zone database is the upper limit for the defined configuration, and it is determined by the amount of flash memory available for storing the defined configuration.

Table 40 Zoning database limitations

Fabric OS version	Maximum database size (KB)
2.4.0	64
2.5.0	64
2.6.0	96
3.0.0	128
3.1.0	96
3.2.0	256
4.0.0, 4.1.0, 4.2.0	128
4.4.0	256
5.0.1	256

Adding a new switch or fabric

When a new switch is added to the fabric, it takes on the zone configuration information from the fabric. Use the `cfgActvShow` command to verify that the zoning information is the same on each switch in the fabric.

If you are adding a switch that is already configured for zoning, use the `cfgClear` and `cfgSave` commands (or use `cfgClear` and `cfgDisable` if there is an effective configuration) before connecting it to the zoned fabric.

Adding a new fabric that has no zone configuration information to an existing fabric is very similar to adding a new switch. All switches in the new fabric inherit the zoning configuration data. If a zone configuration is in effect, the same configuration becomes the enabled configuration. The `cfgActvShow` command displays the same information on all switches in the newly formed fabric.


Before the new fabric can merge successfully, it must pass the following criteria:

- Check the following before merging zones, switches or fabrics.
 - **Zoning licenses:** All switches must have a Zoning license enabled.
 - **Native operating mode:** All switches must be in the native operating mode.
 - **Secure Fabric OS:** If one switch has Secure Fabric OS enabled, all switches in the fabric must have Secure Fabric OS. See the *HP StorageWorks Secure Fabric OS administrator guide* for more information.
- **Merging and segmentation:** The fabric is checked for segmentation during power-up, when a switch is disabled or enabled, or when a new switch is added.

The database is the zone configuration database. This is the data displayed as the `defined` configuration in the `cfgShow` command. It is stored in nonvolatile memory by the `cfgSave` command. This database is a replicated database, which means that all switches in the fabric have a

copy of this database. When a change is made to the defined configuration, the switch where the changes were made must close its transaction for the change to get propagated throughout the fabric.

- **Merging rules:** Observe the following rules when merging zones:
 - Local and adjacent configurations: If the local and adjacent zone database configurations are the same, they remain unchanged after the merge.
 - Effective configurations: If there is an effective configuration between two switches, the zone configurations in effect match.
 - Zone object naming: If a zoning object has the same name in both the local and adjacent defined configurations, the object types and member lists must match. When comparing member lists, the content and order of the members are important.
 - Objects in adjacent configurations: If a zoning object appears in an adjacent defined configuration, but not in the local defined configuration, the zoning object is added to the local defined configuration. The modified zone database must fit in the nonvolatile memory area allotted for the zone database.
 - Local configuration modification: If a local defined configuration is modified because of a merge, the new zone database is propagated to other the switches within the merge request.
- **Merging two fabrics:** Both fabrics have identical zones and configurations enabled. The two fabrics join to make one larger fabric with the same zone configuration across the newly created fabric. If the two fabrics have different zoning configurations, they are merged. If the two fabrics cannot join, the ISL between the switches are segmented.
- **Merge conflicts:** When a merge conflict is present, a merge does not take place and the ISL segments. Use the `switchShow` or `errLogShow` command to obtain additional information about possible merge conflicts, because many non-zone-related configuration parameters can cause conflicts. If the fabrics have different zone configuration data, the system attempts to merge the two sets of zone configuration data. If the zones cannot merge, the ISL is segmented.
A merge is not possible if any of the following conditions exists:
 - Configuration mismatch: Zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
 - Type mismatch: The name of a zone object in one fabric is used for a different type of zone object in the other fabric.
 - Content mismatch: The definition of a zone object in one fabric is different from the definition of zone object with the same name in the other fabric.

 **NOTE:** If the zoneset members on two switches are not listed in the same order, the configuration is considered a mismatch, resulting in the switches being segmented from the fabric. For example, `cfg1 = z1; z2` is different from `cfg1 = z2; z1`, even though members of the configuration are the same. If zoneset members on two switches have the same names defined in the configuration, make sure zoneset members are listed in the same order.

Splitting a fabric

If the connections between two fabrics are no longer available, the fabric segments into two separate fabrics. Each new fabric retains the same zone configuration.

If the connections between two fabrics are replaced and no changes have been made to the zone configuration in either of the two fabrics, the two fabrics merge back into one single fabric. If any changes that cause a conflict have been made to either zone configuration, the fabrics might segment.

Using zoning to administer security


Zones provide controlled access to fabric segments and establish barriers between operating environments. They isolate systems with different uses, protecting individual systems in a heterogeneous environment; for example, when zoning is in secure mode, no merge operations occur.

HP Advanced Zoning is configured on the primary FCS. The primary FCS switch makes zoning changes and other security-related changes. The primary FCS switch also distributes zoning to all other switches in the secure fabric. All existing interfaces can be used to administer zoning, depending on the policies.

You must perform zone management operations from the primary FCS switch using a zone management interface, such as telnet or Advanced Web Tools. You can alter a zoning database, provided you are connected to the primary FCS switch.

When two secure fabrics join, the traditional zoning merge does not occur. Instead, a zoning database is downloaded from the primary FCS switch of the merged secure fabric. When E_Ports are active between two switches, the name of the FCS server and a zoning policy set version identifier are exchanged between the switches. If the views of the two secure fabrics are the same, the fabric's primary FCS server downloads the zoning database and security policy sets to each switch in the fabric. If there is a view conflict, the E_Ports are segmented due to incompatible security data.

As part of zoning architecture, you must determine which of the two basic zoning architectures (hard or soft) works best for your fabric. With time and planning, the basic hard zone configuration works for most sites. If a site has additional security needs, use the additional layer of Secure Fabric OS, apart from the standard zoning architecture.

 **NOTE:** Secure Fabric OS requires the activation of an HP security license and an Advanced Zoning license.

Resolving zone conflicts

Zone conflicts can be resolved by saving a configuration file with the `configUpload` command, examining the zoning information in the file, and performing a cut and paste operation so that the configuration information matches in the fabrics being merged.

After examining the configuration file, you can choose to resolve zone conflicts by using the `cfgClear` command, followed by the `cfgDisable` command on the incorrectly configured segmented fabric, followed by a `portDisable` or `portEnable` command on one of the ISL ports that connects the fabrics. This causes a merge, making the fabric consistent with the correct configuration.


 **CAUTION:** Be careful using the `cfgClear` command; it deletes the defined configuration.

Table 41 lists considerations for zoning architecture.

Table 41 Considerations for zoning architecture

Item	Description
Type of zoning: hard or soft (session-based)	If security is a priority, HP recommends hard zoning.
Use of aliases	The use of aliases is optional with zoning. Using aliases requires structure when defining zones. Aliases aid administrators of zoned fabric in understanding the structure and context.
Security requirements	Evaluate the security requirements of the fabric. If additional security is required, add HP Secure Fabric OS into the fabric.
Interoperability Fabric	If the fabric includes a third-party switch product, only WWN zoning is supported. Other types of zoning, including QuickLoop, are not supported.
QLFA zones	Evaluate if the fabric will have QLFA or QuickLoop in it, and consider the following items before creating and setting up QLFA zones: QuickLoop Zoning—QuickLoop and QuickLoop zones cannot run on Fabric OS 4.1.0 or later. However, Fabric OS can manage (create, remove, update) QuickLoop zones. QLFA—A switch running Fabric OS 4.1.0 or later cannot have a Fabric Assist host directly connected to it. However, such a switch can be part of a Fabric Assist zone if a Fabric Assist host is connected to a compatible switch in the fabric.
Testing	Testing a (new) zone configuration. Before implementing a zone, the user should run the Zone Analyzer from Advanced Web Tools to isolate any possible problems. This is especially useful as fabrics increase in size.
Effect of changes in a production fabric	Zone changes in a production fabric can result in a disruption of I/O under conditions where an RSCN is issued as a result of a zone change, and the HBA is unable to process the RSCN fast enough. Though RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, HP recommends performing zone changes only when the resulting behavior is predictable and acceptable. Changing HBA drivers can rectify the situation.
Confirming operation	After changing or enabling a zone configuration, confirm that the nodes and storage are able to identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.

14 Administering advanced performance monitoring

This topic contains procedures for the HP Advanced Performance Monitoring licensed feature:

Based on HP Frame Filtering technology and a unique performance counter engine, advanced performance monitoring is a comprehensive tool for monitoring the performance of networked storage resources. It supports direct-attach, loop, and switched fabric Fibre Channel SAN topologies by:

- Monitoring transaction performance from source to destination
- Reporting CRC error measurement statistics
- Measuring HP ISL Trunking performance and resource usage

Further features are provided through Web Tools:

- Measuring device performance by port, AL_PA, and LUN
- Comparing IP versus SCSI traffic on each port
- Providing a library of predefined graphs


 **NOTE:** The 4/8 SAN Switch and 4/16 SAN Switch running Fabric OS 5.x do not display AL_PA measurements for end-to-end (EE) monitors. They provide port CRC reports through Advanced Web Tools. The SAN Switch 4/32 switch running Fabric OS 4.4.0 does not display AL_PA measurements for EE monitors. It provides port CRC reports through Advanced Web Tools.


Table 42 lists commands associated with advanced performance monitoring. For detailed information on these commands, see the *HP StorageWorks Fabric OS 5.x command reference guide*.

Table 42 Advanced performance monitoring commands

Command	Description
perfAddEEMonitor	Add an EE monitor to a port.
perfAddIPMonitor	Add an IP monitor to a port.
perfAddReadMonitor	Add a SCSI Read monitor to a port.
perfAddRwMonitor	Add a SCSI Read and Write monitor to a port.
perfAddSCSIMonitor	Add a SCSI traffic frame monitor to a port.
perfAddUserMonitor	Add a filter-based monitor to a port.
perfAddWriteMonitor	Add a SCSI Write monitor to a port.
perfCfgClear	Clear the performance monitoring settings from nonvolatile (flash) memory.
perfCfgRestore	Restore performance monitoring settings from nonvolatile (flash) memory.
perfCfgSave	Save the current performance monitoring settings to nonvolatile (flash) memory.
perfClrAlpaCrc	Clear an AL_PA device CRC count by the port and AL_PA.
perfClearEEMonitor	Clear EE monitor counters on a port.
perfClearFilterMonitor	Clears filter-based monitor counters.
perfDelEEMonitor	Delete an EE monitor on port.
perfDelFilterMonitor	Delete a filter-based monitor.
perfMonitorClear	Clear statistics counters of EE, filter-based, and ISL monitors on a port.
perfMonitorShow	Display EE, filter-based, and ISL monitors on a port.

Table 42 Advanced performance monitoring commands (continued)

Command	Description
<code>perfSetPortEEMask</code>	Set the overall mask for EE monitors.
<code>perfShowAlpaCrc</code>	Display the AL_PA CRC count by port or by AL_PA.
<code>perfShowEEMonitor</code>	Show user-defined EE monitors.
<code>perfShowFilterMonitor</code>	Show filter-based monitors.
<code>perfShowPortEEMask</code>	Display the current EE mask of a port.

 **NOTE:** The command examples in this chapter use the slot/port syntax required by the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director. For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32, use only the port number in the commands.

Displaying and clearing the CRC error count

You can use the `perfShowAlpaCrc` command to display the CRC error count for all AL_PA devices or for a single AL_PA on a specific active L_Port. The following is an example that displays the CRC error count for all AL_PA devices on a port:

```
switch:admin> perfshowalpacrc 1/1
AL_PA      CRC count
-----
0xd9              0
```

The following is an example that displays the CRC error count for a single AL_PA device on a port:

```
switch:admin> perfshowalpacrc 1/1, 0xd9
The CRC count at ALPA 0xd9 on port 1 is 0x000000000.
```

The following is an example that clears the CRC error count:

```
switch:admin> perfclrallpacrc 1/1, 0xd9
CRC error count at AL_PA 0xd9 on port 1 is cleared.
switch:admin> perfclrallpacrc 1/1
No AL_PA value is specified. This will clear all AL_PA CRC
counts on port 1. Do you want to continue? (yes, y, no, n): [no] y
Please wait ...
All alpa CRC counts are cleared on port 1.
```

In Fabric OS 3.1.0, 4.1.0, and later versions, you can use the `portStatsClear` command clears AL_PA- based CRC error counters for all the ports in the same group.

Monitoring EE performance

EE performance monitoring counts the number of words and CRC errors in Fibre Channel frames for a specified SID-DID pair. An EE performance monitor includes the following counts:

- RX_COUNT (words in frames received at the port)
- TX_COUNT (words in frames transmitted from the port)
- CRC_COUNT (frames with CRC errors received at or transmitted from the port)


To enable EE performance monitoring, you must configure an EE monitor on a port, specifying the SID-DID pair (in hexadecimal). The monitor counts only those frames with a matching SID and DID.

Each SID or DID has three fields, listed in the following order:

- Domain ID (DD)
- Area ID (AA)
- AL_PA (PP)

For example, the SID 0x118a0f denotes DD 0x11, AA 0x8a, and AL_PA 0x0f.

You can monitor EE performance using the `perfMonitorShow` command, as described in “[Displaying monitor counters](#)” on page 207. You can clear EE counters using the `perfMonitorClear` command, as described in “[Clearing monitor counters](#)” on page 209.

 **NOTE:** For EE monitors, CRC counters are not displayed on the 4/8 SAN Switch, 4/16 SAN Switch, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, nor on the SAN Switch 4/32, nor on the 4/256 SAN Director.

Adding EE monitors

An EE monitor counts the following items for a port: number of words received, number of words transmitted, and number of CRC errors detected in frames.

The 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and Brocade 4Gb SAN Switch for HP p-Class BladeSystem, Core Switch 2/64, and SAN Director 2/128 allow up to eight EE monitors.

The SAN Switch 4/32 and 4/256 SAN Director allow up to 256 EE monitors shared by all ports. (The number of ISLs configured on the switch affects the amount of resources available for EE monitors.)

EE monitors cannot be added to ISLs.

The monitor count is qualified using either of following conditions:

- For frames received at the port with the EE monitor installed, the frame SID is the same as `SourceID` and the frame DID is the same as `DestID`. The RX_COUNT and CRC_COUNT are updated accordingly.
- For frames transmitted from the port with the EE monitor installed, the frame DID is the same as `SourceID` and the frame SID is the same as `DestID`. The TX_COUNT and CRC_COUNT are updated accordingly.


 **NOTE:** How the area ID for a port relates to the port number depends upon the PID format used by the fabric. See “[Configuring the PID format](#)” on page 213 for more information.

Figure 10 shows two devices:

- Host A is connected to domain 5 (0x05), switch area ID 18 (0x12), AL_PA 0x00 on Switch X.
- Dev B is a storage device connected to domain 17 (0x11), switch area ID 30 (0x1e), AL_PA 0xef on Switch Y.

NOTE: EE performance monitoring looks at traffic on the receiving port respective to the SID only. In Figure 10, if you add a monitor to slot 2, port 2, on Switch X, specifying Dev B as the SID and Host A as the DID, no counters (except CRC) are incremented.

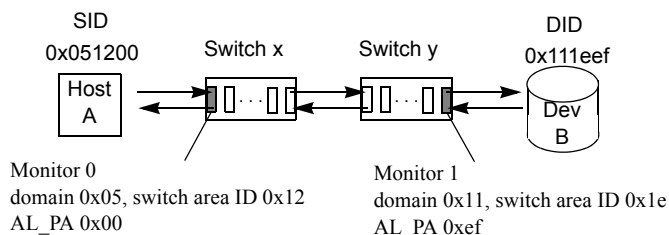


Figure 10 Setting EE monitors on a port

Monitoring the traffic from Host A to Dev B

Add Monitor 0 to slot 2, port 2 on Switch X, specifying 0x051200 as the SID and 0x111eef as the DID, as shown in the following example:

```
switch:admin> perfaddeemonitor 2/2, "0x051200" "0x111eef"  
End-to-End monitor number 0 added.
```

Monitor 0 counts the frames that have an SID of 0x051200 and a DID of 0x111eef. For monitor 0, RX_COUNT is the number of words from Host A to Dev B, TX_COUNT is the number of words from Dev B to Host A, and CRC_COUNT is the number of frames in both directions with CRC errors.

Monitoring the traffic from Dev B to Host A

Add Monitor 1 to slot 2, port 14 on Switch y, specifying 0x111eef as the SID and 0x051200 as the DID, as shown in the following example:

```
switch:admin> perfaddeemonitor 2/14, "0x111eef" "0x051200"  
End-to-End monitor number 1 added.
```

Monitor 1 counts the frames that have an SID of 0x111eef and a DID of 0x051200. For monitor 1, RX_COUNT is the number of words from Dev B to Host A, TX_COUNT is the number of words from Host A to Dev B, and CRC_COUNT is the number of frames in both directions with CRC errors.

Figure 11 shows several switches and the correct ports on which to add performance monitors for a specified SID-DID pair.

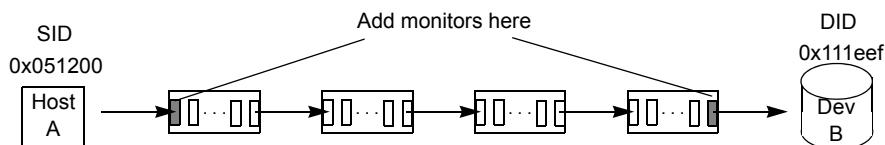


Figure 11 Proper placement of EE performance monitors

Setting a mask for EE monitors

EE monitors count the number of words in Fibre Channel frames that match a specific SID-DID pair. If you want to match only part of the SID or DID, you can set a mask on the port to compare only certain parts of the SID or DID. By default, the frame must match the entire SID and DID to trigger the monitor. By setting a mask, you can choose to have the frame match only one or two of the three fields (Domain ID, Area ID, and AL_PA) to trigger the monitor.

△ **CAUTION:** Only one mask per port can be set. When you set a mask, all existing EE monitors are deleted.

You can specify a mask using the `perfSetPortEeMask` command in the form `dd:aa:pp`, where `dd` is the domain ID mask, `aa` is the area ID mask, and `pp` is the AL_PA mask. The values for `dd`, `aa`, and `pp` are either `ff` (the field must match) or `00` (the field is ignored). The default EE mask value is `ff:ff:ff`. The command sets the mask for all EE monitors of a port. If any EE monitors are programmed on a port when the `perfSetPortEeMask` command is issued, a message is displayed as shown in the following example:

```
switch:admin> perfsetporteemask 1/2, "00:00:ff"
EE monitors are currently programmed on this port. Changing EE mask
for this port will cause ALL EE monitors on this port to be deleted.
Do you want to continue? (yes, y, no, n): [no] y

EE mask on port <port-number> is set and EE monitors were deleted
```

The `perfSetPortEeMask` command sets a mask for the Domain ID, Area ID, and AL_PA of the SIDs and DIDs for frames transmitted from and received by the port.

Figure 12 shows the mask positions in the command. A mask (`ff`) is set on slot 1, port 2 to compare the AL_PA fields on the SID and DID in all frames (transmitted and received) on port 2. The frame SID and DID must match only the AL_PA portion of the specified SID-DID pair. Each port can have only one EE mask. The mask is applied to all EE monitors on the port. Individual masks for each monitor on the port cannot be specified.

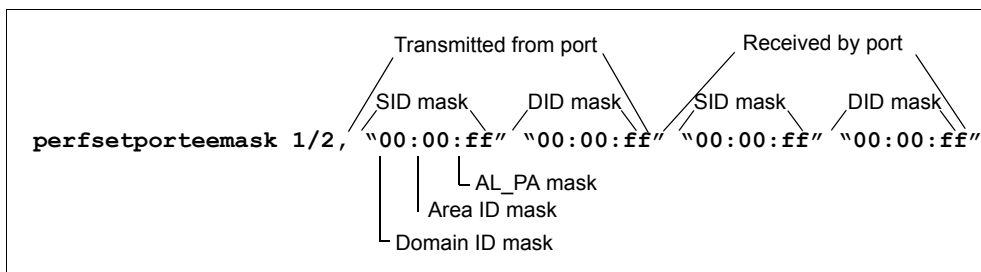


Figure 12 Mask positions for EE monitors

Displaying the current EE mask of a port

Issue the `perfShowPortEeMask` command.

Adding standard filter-based monitors

Table 43 lists the commands for adding standard filter-based monitors to a port.

Table 43 Commands to add filter-based monitors

Telnet command	Description
perfAddReadMonitor	Counts the number of SCSI read commands.
perfAddWriteMonitor	Counts the number of SCSI write commands.
perfAddRwMonitor	Counts the number of SCSI read and write commands.
perfAddScsiMonitor	Counts the number of SCSI traffic frames.
perfAddIpMonitor	Counts the number of IP traffic frames.

The following example adds filter-based monitors to slot 1, port 2 and displays the results:

```
switch:admin> perfaddreadmonitor 1/2
SCSI Read filter monitor #0 added
switch:admin> perfaddwritemonitor 1/2
SCSI Write filter monitor #1 added
switch:admin> perfaddrwmonitor 1/2
SCSI Read/Write filter monitor #2 added
switch:admin> perfaddscsimonitor 1/2
SCSI traffic frame monitor #3 added
switch:admin> perfaddipmonitor 1/2
IP traffic frame monitor #4 added
switch:admin> perfmonitorshow --class FLT 1/2
There are 5 filter-based monitors defined on port 2.
```

KEY	ALIAS	OWNER_APP	OWNER_IP_ADDR	FRAME_COUNT
0	SCSI Read	TELNET	N/A	0x0000000000000000
1	SCSI Write	TELNET	N/A	0x0000000000000000
2	SCSI R/W	TELNET	N/A	0x0000000000000000
3	SCSI Frame	TELNET	N/A	0x0000000000000000
4	IP Frame	TELNET	N/A	0x0000000000000000

Adding custom filter-based monitors

In addition to the standard filters—read, write, read/write, SCSI frame and IP frame—you can create custom filters to gather statistics that fit your needs.

To define a custom filter, use the `perfAddUserMonitor` command. With this command, you must specify a series of offsets, masks, and values. For all transmitted frames, the switch performs these tasks:

- Locates the byte found in the frame at the specified offset
- Applies the mask to the byte found in the frame
- Compares the value with the given values in the `perfAddUserMonitor` command
- Increments the filter counter if a match is found

The following number of offsets can be specified:

- SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, and SAN Director 2/128 (Fabric OS 4.0.0 or later): Up to two different offsets per port (one offset when FMS is enabled)
- 4/256 SAN Director (Fabric OS 5.0.1 or later): Up to 15 different offsets per port (14 offsets when `fmsmode` is enabled)
- SAN Switch 2/8-EL and SAN Switch 2/16 (Fabric OS 3.0.0 or later): Up to three different offsets per port

- SAN Switch 4/32 (Fabric OS 4.4.0 or later): Up to 15 different offsets per port (14 offsets when FMS is enabled)
- 4/8 SAN Switch, 4/16 SAN Switch, and Brocade 4Gb SAN Switch for HP p-Class BladeSystem (Fabric OS 5.0.1): Up to 7 different offsets per port (6 offsets when fmsmode is enabled)

You can specify up to four values to compare against each offset. If more than one offset is required to properly define a filter, the bytes found at each offset must match one of the given values for the filter to increment its counter. If one or more of the given offsets does not match any of the given values, the counter does not increment.

The value of the offset must be between 0 and 63, in decimal format. Byte 0 indicates the first byte of the Start of Frame (SOF), byte 4 is the first byte of the frame header, and byte 28 is the first byte of the payload. Thus, only the SOF, frame header, and first 36 bytes of payload can be selected as part of a filter definition. Offset 0 is a special case, which can be used to monitor the first 4 bytes of the frame. When the offset is set to 0, the values 0–7 that are checked against that offset are predefined, as shown in [Table 44](#).

Table 44 Predefined values at offset 0

Value	SOF	Value	SOF
0	SOFF	4	SOFi2
1	SOFc1	5	SOFn2
2	SOFi1	6	SOFi3
3	SOFn1	7	SOFn3

If the switch does not have sufficient resources to create a given filter, other filters might have to be deleted to free resources.

Adding filter-based monitors

```
switch:admin> perfaddusermonitor 4/2, "12, 0xff, 0x05, 0x08; 9, 0xff, 0x02" "FCP/IP"
User monitor #5 added
switch:admin> perfaddusermonitor 1/2, "0, 0xff, 6"
User Monitor #6 added
```

In the preceding example, two filter-based monitors are added. The first monitor (#5) counts all FCP and IP frames transmitted from domain 0x02 for slot 4, port 2. The FCP and IP protocols are selected by monitoring offset 12, mask 0xff and matching values of 0x05 or 0x08. Domain 2 is selected by monitoring offset 9, mask 0xff, and matching a value of 0x02. The monitor counter is incremented for all outgoing frames from port 2 where byte 9 is 0x02 and byte 12 is 0x05 or 0x08.

The second monitor (#6) is for SOFi3 on slot 1, port 2.

Deleting filter-based monitors

1. List the valid monitor numbers by entering the `perfShowFilterMonitor` command.
2. Issue the `perfDelFilterMonitor` command to delete a specific monitor. If you do not specify which monitor number to delete, you are asked whether you want to delete all entries.

The following example displays the monitors on slot 1, port 4 using the `perfShowFilterMonitor` command (the monitor numbers are listed in the `KEY` column) and deletes monitor number 1 on slot 1:

```
switch:admin> perfshowfiltermonitor 1/4
There are 4 filter-based monitors defined on port 4.
KEY    ALIAS    OWNER_APP    OWNER_IP_ADDR    FRAME_COUNT
-----
0  SCSI Read  TELNET                N/A              0x00000000000002208
1  SCSI Write TELNET                N/A              0x0000000000000464a
2  SCSI R/W   TELNET                N/A              0x0000000000000fd8c
3  SCSI Frame WEB_TOOLS            192.168.169.40   0x0000000000002c229
switch:admin> perfdelfiltermonitor 1/4, 1
The specified filter-based monitor is deleted.
```

Monitoring ISL performance

ISL monitoring is set up on E_Ports in release 4.4.0 and later.

An ISL monitor measures traffic to all reachable destination domains for an ISL, showing which destination domain is consuming the most traffic. If there are more than 16 domains, the monitor samples traffic and extrapolates the measurement.

You can monitor ISL performance using the `perfMonitorShow` command, as described in “[Displaying monitor counters](#)” on page 207.” You can clear ISL counters using the `perfMonitorClear` command, as described in “[Clearing monitor counters](#)” on page 209.

Monitoring trunks

For trunked ISLs on Fabric OS 4.x switches, monitoring is set only on the master ISL, which communicates with the associated slave ISLs. For Fabric OS 3.x switches, monitoring can be set on slave ISLs.


EE monitors are not supported for ISLs.

The SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director support eight filter-based monitors for trunks.

The 4/8 SAN Switch, 4/16 SAN Switch, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32 support 12 filter-based monitors for trunks.

Displaying monitor counters

Use the `perfMonitorShow` command to display the monitors on a specified port. For EE counters, you can display either the cumulative count of the traffic detected by the monitors or a snapshot of the traffic at specified intervals.

 **NOTE:** The 4/8 SAN Switch, 4/16 SAN Switch, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32 outputs do not include CRC counts.

The command format is:

```
perfmonitorshow --class monitor_class [slotnumber/]portnumber [interval]
```

where:

monitor_class Specifies the monitor class, which can be EE, FLT (filter-based), or ISL. The --class monitor_class operand is required.

slotnumber Specifies the slot number for a Core Switch 2/64, SAN Director 2/128, or 4/256 SAN Director. For all other switches, this operand is not required. The slot number must be followed by a slash (/) and the port number, so that each port is represented by both slot number (1 through 4 or 7 through 10) and port number (0 through 15).

The HP StorageWorks director has a total of 10 slots. Slot numbers 5 and 6 are CP blades; slots 1 through 4 and 7 through 10 are port blades. For 16-port blades, there are 16 ports, counted from the bottom, numbered 0 to 15. For 32-port blades, there are 32 ports numbered 0 to 31.

portnumber Specifies a port number. Valid values for port number vary, depending on the switch type. This operand is required.

interval Specifies an interval in seconds. The interval must be greater than or equal to five seconds. For EE monitoring, the Tx and Rx counts are measured in bytes. This operand is optional.

The following example displays an EE monitor on a port at an interval of every 6 seconds:

```
switch:admin> perfMonitorShow --class EE 4/5 6
perfmonitorshow 53, 6: Tx/Rx are # of bytes and crc is # of crc errors
      0          1          2          3          4
-----
crc   Tx   Rx  crc   Tx   Rx  crc   Tx   Rx  crc   Tx   Rx  crc   Tx   Rx
=====
0     0     0   0     0     0   0     0     0   0     0     0   0     0     0
0   53m   4.9m 0   53m   4.9m 0   53m   4.9m 0   53m   4.9m 0   53m   4.9m 0
0   53m   4.4m 0   53m   4.4m 0   53m   4.4m 0   53m   4.4m 0   53m   4.4m 0
0   53m   4.8m 0   53m   4.8m 0   53m   4.8m 0   53m   4.8m 0   53m   4.8m 0
0   53m   4.6m 0   53m   4.6m 0   53m   4.6m 0   53m   4.6m 0   53m   4.6m 0
0   53m   5.0m 0   53m   5.0m 0   53m   5.0m 0   53m   5.0m 0   53m   5.0m 0
0   53m   4.8m 0   53m   4.8m 0   53m   4.8m 0   53m   4.8m 0   53m   4.8m 0
0   53m   4.5m 0   53m   4.5m 0   53m   4.5m 0   53m   4.5m 0   53m   4.5m 0
0   52m   4.5m 0   52m   4.5m 0   52m   4.5m 0   52m   4.5m 0   52m   4.5m 0
0   52m   5.0m 0   52m   5.0m 0   52m   5.0m 0   52m   5.0m 0   52m   5.0m 0
0   52m   4.5m 0   52m   4.5m 0   52m   4.5m 0   52m   4.5m 0   52m   4.5m 0
0   52m   4.6m 0   52m   4.6m 0   52m   4.6m 0   52m   4.6m 0   52m   4.6m 0
```

The following example displays EE monitors on a port:

```
switch:admin> perfMonitorShow --class EE 4/5
There are 7 end-to-end monitor(s) defined on port 53.
KEY      SID      DID      OWNER_APP      OWNER_IP_ADDR      TX_COUNT      RX_COUNT      CRC_COUNT
-----
0  0x58e0f  0x1182ef  TELNET          N/A              0x0000000000000000  0x0000000000000000  0x0000000000000000
0  0x21300  0x21dda  TELNET          N/A              0x00000004d0ba9915  0x0000000067229e65  0x0000000000000000
1  0x21300  0x21ddc  TELNET          N/A              0x00000004d0baa754  0x0000000067229e65  0x0000000000000000
2  0x21300  0x21de0  TELNET          N/A              0x00000004d0bab3a5  0x0000000067229e87  0x0000000000000000
3  0x21300  0x21de1  TELNET          N/A              0x00000004d0bac1e4  0x0000000067229e87  0x0000000000000000
4  0x21300  0x21de2  TELNET          N/A              0x00000004d0bad086  0x0000000067229e87  0x0000000000000000
5  0x11000  0x21fd6  WEB_TOOLS      192.168.169.40    0x00000004d0bade54  0x0000000067229e87  0x0000000000000000
6  0x11000  0x21fe0  WEB_TOOLS      192.168.169.40    0x00000004d0baed41  0x0000000067229e98  0x0000000000000000
```

The following example displays filter-based monitor on a port at an interval of every 6 seconds:

```
switch:admin> perfMonitorShow --class FLT 2/5 6
perfmonitorshow 21, 6
```

0	1	2	3	4	5	6
#Frames	#Frames	#Frames	#Frames	#Frames	#Frames	#Frames
0	0	0	0	0	0	0
26k	187	681	682	682	494	187
26k	177	711	710	710	534	176
26k	184	734	734	734	550	184
26k	182	649	649	649	467	182
26k	188	754	755	755	567	184
26k	183	716	716	717	534	183
26k	167	657	656	655	488	167
26k	179	749	749	749	570	179
26k	164	752	752	752	588	164
26k	190	700	700	700	510	190
26k	181	701	701	701	520	181
26k	200	750	750	751	550	201
26k	180	692	692	691	512	179
26k	179	696	696	696	517	179
26k	187	720	720	720	533	187
26k	200	722	722	722	522	200
26k	204	717	717	717	513	204

The following example displays filter monitor information on a port:

```
switch:admin> perfMonitorShow --class FLT 2/5
There are 7 filter-based monitors defined on port 21.
```

KEY	ALIAS	OWNER_APP	OWNER_IP_ADDR	FRAME_COUNT
0	SCSI_Frame	TELNET	N/A	0x000000000002c2229
1	SCSI_WR	TELNET	N/A	0x0000000000000464a
2	SCSI_RW	TELNET	N/A	0x0000000000000fd8c
3	SCSI_RW	WEB_TOOLS	192.168.169.40	0x00000000000007ba3
4	SCSI_RW	WEB_TOOLS	192.168.169.190	0x00000000000004f0e
5	SCSI_RD	WEB_TOOLS	192.168.169.40	0x00000000000002208
6	SCSI_WR	WEB_TOOLS	192.168.169.40	0x0000000000000033a

The following example displays ISL monitor information on a port:

```
switch:admin> perfMonitorShow --class ISL 1/1
Total transmit count for this ISL: 1462326
Number of destination domains monitored: 3
Number of ports in this ISL: 2
Domain 97: 110379 Domain 98: 13965
Domain 99: 1337982
```

Clearing monitor counters


Before you clear statistics counters, verify the valid monitor numbers on a specific port using the `perfMonitorShow` command, to make sure the correct monitor counters are cleared. To clear statistics counters for all monitors or a specified monitor, use the `perfMonitorClear` command. After the command has been executed, the telnet shell confirms that the counters on the monitor have been cleared.

The command format is:

```
perfmonitorclear --class monitor_class [slotnumber/]portnumber [monitorId]
```

where:

<i>monitor_class</i>	Specifies the monitor class, which can be one of EE, FLT (filter-based), or ISL. The <code>--class <i>monitor_class</i></code> operand is required.
<i>slotnumber</i>	<p>Specifies the slot number for a Core Switch 2/64, SAN Director 2/128, or 4/256 SAN Director. For all other switches, this operand is not required. The slot number must be followed by a slash (/) and the port number, so that each port is represented by both slot number (1 through 4 or 7 through 10) and port number (0 through 15).</p> <p>The HP StorageWorks director has a total of 10 slots. Slot numbers 5 and 6 are CP blades; slots 1 through 4 and 7 through 10 are port blades. For 16-port blades, there are 16 ports, counted from the bottom, numbered 0 to 15. For 32-port blades, there are 32 ports numbered 0 to 31.</p>
<i>portnumber</i>	Specifies a port number. Valid values for port number vary, depending on the switch type. This operand is required.
<i>monitorId</i>	Specifies the monitor number to clear. Monitor numbers are defined when you create the monitor on a port. This operand is optional. If not specified, all monitor counters on the port are cleared. This operand does not apply to ISL monitors.

 **NOTE:** In Fabric OS 3.1.0 and 4.1.0 (or later) the `portStatsClear` command clears AL_PA-based CRC error counters for all the ports in the same group.

The following example clears statistics counters for an EE monitor:

```
switch:admin> perfMonitorClear --class EE 1/2 5
End-to-End monitor number 5 counters are cleared

switch:admin> perfMonitorClear --class EE 1/2
This will clear ALL EE monitors' counters on port 2, continue?
(yes, y, no, n): [no] y
```

The following examples clears statistics counters for a filter-based monitor:

```
switch:admin> perfMonitorClear --class FLT 1/2 4
Filter-based monitor number 4 counters are cleared

switch:admin> perfMonitorClear --class FLT 1/2
This will clear ALL filter-based monitors' counters on port 2, continue?
(yes, y, no, n): [no] y
```

The following example clears statistics counters for an ISL monitor:

```
switch:admin> perfMonitorClear --class ISL 1
This will clear ISL monitor on port 1, continue? (yes, y, no, n): [no] y
```

Saving and restoring monitor configurations

To save the current EE and filter monitor configuration settings into nonvolatile memory, use the `perfCfgSave` command. For example:

```
switch:admin> perfCfgsave  
This will overwrite previously saved Performance Monitoring settings in  
FLASH ROM. Do you want to continue? (yes, y, no, n): [no] y  
Please wait... Committing configuration...done.  
Performance monitoring configuration saved in FLASH ROM.
```

To restore a saved monitor configuration, use the `perfCfgRestore` command. For example, to restore the original performance monitor configuration after making several changes:

```
switch:admin> perfCfgrestore  
This will overwrite current Performance Monitoring settings in RAM. Do you  
want to continue? (yes, y, no, n): [no] y  
Please wait... Performance monitoring configuration restored from FLASH  
ROM.
```

To clear the previously saved performance monitoring configuration settings from nonvolatile memory, use the `perfCfgClear` command. For example:

```
switch:admin> perfCfgclear  
This will clear Performance Monitoring settings in FLASH ROM. The RAM  
settings won't change. Do you want to continue? (yes, y, no, n): [no] y  
Please wait... Committing configuration...done.  
Performance Monitoring configuration cleared from FLASH.
```

Collecting performance data

Data collected through advanced performance monitoring is deleted when the switch is rebooted. Using the HP Fabric Manager software application version 4.4.0 (or later), you can store performance data persistently. For details on this feature, see the *HP StorageWorks Fabric Manager 5.x administrator guide*.

A Configuring the PID format

Port identifiers (PIDs) are used by the routing and zoning services in Fibre Channel fabrics to identify ports in the network. All devices in a fabric must use the same PID format, so when you add new equipment to your SAN, you might need to change the PID format on legacy equipment.

About PIDs and PID binding

The PID is a 24-bit address built from three 8-bit fields:

- domain
- area_ID
- AL_PA

Many scenarios cause a device to receive a new PID, for example, unplugging the device from one port and plugging it into a different port as part of fabric maintenance, or changing the domain ID of a switch, which might be necessary when merging fabrics, or changing compatibility mode settings.

Some device drivers use the PID to map logical disk drives to physical Fibre Channel counterparts. Most drivers can either change PID mappings dynamically (called *dynamic PID binding*) or use the WWN of the Fibre Channel disk for mapping (called *WWN binding*).

Some older device drivers behave as if a PID uniquely identifies a device (they use *static PID binding*). These device drivers should be updated, if possible, to use WWN or dynamic PID binding instead, because static PID binding creates problems in many routine maintenance scenarios. Fortunately, very few device drivers still behave this way. Many current device drivers enable you to select static PID binding as well as WWN binding. Select static binding only if there is a compelling reason, and only after you have evaluated the impact of doing so.

Summary of PID formats

HP StorageWorks switches employ these types of PID formats:

- VC encoded: The format defined by the Fibre Channel Storage Switch 8 and Fibre Channel Storage Switch 16. Connections to these switches are not supported in Fabric OS 4.0.0 and later.
- Native: Introduced with the StorageWorks 1 GB switches, this format supports up to 16 ports per switch.
- Core: The default for the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, SAN Switch 4/32, Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director and is the recommended format for HP StorageWorks switches and fabrics. It uses the entire 8-bit address space and directly uses the port number as the area_ID. It supports up to 256 ports per switch.
- Extended edge: This format generates the same PID for a port on switches with 16 or fewer ports as would native PID format, but it also supports up to 256 ports per domain. It should be used only in cases where you cannot upgrade devices to dynamic PID binding and you absolutely cannot reboot your servers.

Extended edge PID is supported in Fabric OS 2.6.2 and later, 3.1.2 and later, and 4.2.0 and later.

In addition to the PID formats list here, interoperability mode supports additional PID formats that are not discussed in this guide.

Impact of changing the fabric PID format

If your fabric contains switches that use Native PID, HP recommends that you change the format to Core PID before you add the new, higher port count switches and directors. HP recommends that you use Core PID when upgrading the Fabric OS version on HP StorageWorks 1 GB and 2 GB switches.

Depending on your situation, the PID change might or might not entail fabric downtime:

- If you are running dual-fabrics with multipathing software, you can update one fabric at a time without disrupting traffic. Move all traffic onto one fabric in the SAN and update the other fabric, and then move the traffic onto the updated fabric, and update the final fabric.
- Without dual-fabrics, HP recommends stopping traffic. This is the case for many routine maintenance situations, so dual-fabrics are always recommended for uptime-sensitive environments. If your fabric contains devices that employ static PID binding, or you do not have dual-fabrics, you must schedule downtime for the SAN to change the PID format.

You can find details on the impact of PID changes in the following publications, which are available on the HP web site <http://welcome.hp.com/country/us/en/prodserv/storage.html>.

The following sections describe various impacts of PID format changes in greater detail.

Host reboots

In some Fibre Channel SAN environments, storage devices and host servers are bound to the host operating system by their PIDs (called their *Fibre Channel addresses*). In these environments, the hosts and target HBAs in a SAN need to know the full 24-bit PIDs of the hosts and targets they are communicating with, but they do not care how the PIDs are determined. If a storage device PID is changed, however, the host must reestablish a new binding, which requires the host to be rebooted.

With the introduction of the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32 and the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, the Native PID format used in earlier switches was supplemented with the Core PID format, which is capable of addressing higher port counts. Changing from Native PID format to Core PID format changes the PID, which requires hosts that use port binding to be rebooted.

Static PID mapping errors

If possible, do not use drivers that employ static PID binding.

With the WWN or dynamic PID binding most typically used with drivers, changing the device's PID does not affect the PID mapping. However, before updating the PID format, it is necessary to determine whether any devices in the SAN use static PID binding.

For those few drivers that do use static PID binding, changing the PID format breaks the mapping, which must be fixed either by rebooting the host or by using a manual update procedure on the host.

To correct mapping errors caused by static PID binding, see the following sections:

- See "[Evaluating the fabric](#)" on page 216 for details on finding devices that use static PID binding, and then see "[Online update](#)" on page 218 or "[Offline update](#)" on page 218 for recommendations.
- See "[Converting port number to area ID](#)" on page 222 for instructions.

Changes to configuration data

[Table 45](#) lists various combinations of before-and-after PID formats, and indicates whether the configuration is affected.

△ **CAUTION:** After changing the fabric PID format, if the change invalidates the configuration data (see [Table 45](#) to determine this), do not download old (pre-PID format change) configuration files to any switch on the fabric.

Table 45 Effects of PID format changes on configurations

PID format before change	PID format after change	Configuration effect
Native	Extended Edge	No impact
Extended Edge	Native	No impact
Native	Core	You must: <ul style="list-style-type: none">• Reenable zoning, if there is an active zone set and it uses port zones.• You do not need to reconfigure DID if using:<ul style="list-style-type: none">• Performance monitoring• The configure command to change the PID format• The performance monitor database• The zoning database• Security dcc database (if secure mode is enabled) The DID is converted by FOS.
Core	Native	
Extended Edge	Core	
Core	Extended Edge	

After changing the fabric PID format and verifying correct fabric operation, resave configuration data by running the `configUpload` command.

Before downgrading firmware, change the PID back to supported PIDs, such as Core PID. If the database is converted, save the converted database, and then download the earlier OS.

Selecting a PID format

All switches in a fabric must use the same PID format, so if you add a switch that uses a different PID format to a fabric, the switch segments from the fabric. The format you select for your fabric depends on the mix of switches in the fabric, and to an extent on the specific releases of Fabric OS in use (for example, Extended Edge PID format is available only in Fabric OS 2.6.2 and later, Fabric OS 3.1.2 and later, and Fabric OS 4.2.0 and later).

If you are building a new fabric with switches running various Fabric OS versions, use Core PID format to simplify port-to-area_ID mapping.

Table 46 shows various combinations of existing fabrics, new switches added to those fabrics, and the recommended PID format for that combination. The criteria for the recommendations are first to eliminate host reboots, and second to minimize the need for a host reboot in the future.

Table 46 PID format recommendations for adding new switches

Existing Fabric OS versions; PID format	Switch to be added	Recommendations (in order of preference)
Version 2.6.2 and later; version 3.1.2 and later; Native PID	2.6.2 and later; 3.1.2 and later	<ul style="list-style-type: none"> • Use Native PID format for new switch. Host reboot is not required. • Convert existing fabric to Core PID format, upgrading the version of Fabric OS, if necessary. Set Core PID format for new switch. Host reboot is required. • If devices are bound statically and it is not possible to reboot, convert existing fabric to Extended Edge PID format, upgrading the version of Fabric OS, if necessary. Use Extended Edge PID format for new switch. Host reboot is not required.
	4.2.0 and later	<ul style="list-style-type: none"> • Convert existing fabric to Core PID format, upgrading the version of Fabric OS, if necessary. Set Core PID format for new switch. Host reboot is required. • If devices are bound statically and it is not possible to reboot, convert existing fabric to Extended Edge PID format, upgrading the version of Fabric OS, if necessary. Use Extended Edge PID format for new switch. Host reboot is not required.
Version 2.6.2 and later; version 3.1.2 and later; version 4.2.0 and later; Core PID	2.6.2 and later; 3.1.2 and later; 4.2.0 and later	Use Core PID for new switch. Host reboot is not required.
Version 2.6.2 and later; version 3.1.2 and later; version 4.2.0 and later; Extended Edge PID	2.6.2 and later; 3.1.2 and later; 4.2.0 and later	Use Extended Edge PID for new switch. Host reboot is not required.

Evaluating the fabric

If there is the possibility that your fabric contains host devices with static PID bindings, evaluate the fabric to:

- Find any devices that bind to PIDs
- Determine how each device driver responds to the PID format change
- Determine how any multipathing software responds to a fabric service interruption

If current details about the SAN are already available, it might be possible to skip the data collection step. If not, it is necessary to collect information about each device in the SAN. Any type of device might be able to bind by PID; each device should be evaluated before attempting an online update. This information has broad applicability, because PID-bound devices are not able to seamlessly perform in many routine maintenance or failure scenarios. To evaluate the fabric:

1. Collect device, software, hardware, and configuration data.

The following is a non-comprehensive list of information to collect:

- HBA driver versions
- Fabric OS versions
- RAID array microcode versions
- SCSI bridge code versions
- JBOD drive firmware versions
- Multipathing software versions
- HBA timeout values
- Multipathing software timeout values
- Kernel timeout values
- Configuration of switch

2. Make a list of manually configurable PID drivers.

Some device drivers do not bind by PID, but allow the operator to manually create a PID binding. For example, persistent binding of PIDs to logical drives might be done in many HBA drivers. Make a list of all devices that are configured this way. If manual PID binding is in use, consider changing to WWN binding.

The following are some of the device types that might be manually configured to bind by PID:

- HBA drivers (persistent binding)
- RAID arrays (LUN access control)
- SCSI bridges (LUN mapping)

3. Analyze data.

After you have determined the code versions of each device on the fabric, the devices must be evaluated to find out whether any bind by PID. It might be easiest to work with the support providers of these devices to get this information. If this is not possible, you might need to perform empirical testing. Binding by PID can create management difficulties in a number of scenarios. HP recommends that you not use drivers that bind by PID. If the current drivers do bind by PID, upgrade to WWN-binding drivers if possible.

The drivers shipping by default with HP/UX and AIX at the time of this writing still bind by PID, and so detailed procedures are provided for these operating systems in this chapter. Similar procedures can be developed for other operating systems that run HBA drivers that bind by PID.

There is no inherent PID binding problem with either AIX or HP/UX. It is the HBA drivers shipping with these operating systems that bind by PID. Both operating systems are expected to release HBA drivers that bind by WWN, and these drivers might already be available through some support channels. Work with the appropriate support provider to find out about driver availability.

It is also important to understand how multipathing software reacts when one of the two fabrics is taken offline. If the time-outs are set correctly, the failover between fabrics should be transparent to the users.

Use the multipathing software to manually fail a path before starting maintenance on that fabric.

4. Perform empirical testing.

Empirical testing might be required for some devices, to determine whether they bind by PID. If you are not sure about a device, work with the support provider to create a test environment.

Create as close a match as practical between the test environment and the production environment, and perform an update using the procedure in ["Online update"](#) on page 218.

Devices that bind by PID are unable to adapt to the new format, and one of three approaches must be taken with them:

- A plan can be created for working around the device driver's limitations in such a way as to allow an online update. See the detailed procedures section for examples of how this could be done.
- The device can be upgraded to drivers that do not bind by PID.
- Downtime can be scheduled to reset the device during the core PID update process, which generally allows the mapping to be rebuilt.

If either of the first two options are used, the procedures should again be validated in the test environment.

Determine the behavior of multipathing software, including but not limited to:

- HBA time-out values
- Multipathing software time-out values
- Kernel time-out values

Planning the update procedure

Whether it is best to perform an offline or online update depends on the uptime requirements of the site.

- An offline update requires that all devices attached to the fabric be offline.
- With careful planning, it should be safe to update the core PID format parameter in a live, production environment. This requires dual fabrics with multipathing software. Avoid running backups during the update process, as tape drives tend to be very sensitive to I/O interruption. The online update process is intended for use only in uptime-critical dual-fabric environments, with multipathing software (high-uptime environments should always use a redundant fabric SAN architecture). Schedule a time for the update when the least critical traffic is running.

All switches running any version of Fabric OS 3.1.2 and later or 4.2.0 and later are shipped with the Core PID format enabled, so it is not necessary to perform the PID format change on these switches.

Migrating from manual PID binding (such as persistent binding on an HBA) to manual WWN binding and upgrading drivers to versions that do not bind by PID can often be done before setting the core PID format. This reduces the number of variables in the update process.

Online update

The following steps are intended to provide SAN administrators a starting point for creating site-specific procedures.

1. Back up all data and verify backups.
2. Verify that the multipathing software can switch over between fabrics seamlessly. If there is doubt, use the software's administrative tools to manually disassociate or mark offline all storage devices on the first fabric to be updated.
3. Verify that I/O continues on the other fabric.
4. Disable all switches in the fabric to be updated, one switch at a time, and verify that I/O continues on the other fabric after each switch is disabled.
5. Change the PID format on each switch in the fabric.
6. Reenable the switches in the updated fabric one at a time.

In a core/edge network, enable the core switches first.

7. After the fabric has reconverged, use the `cfgEnable` command to update zoning.
8. Update the bindings for any devices manually bound by PID.

This might involve changing them to the new PIDs, or preferably changing to WWN binding.

For any devices bound by PID, two options exist:

- Execute a custom procedure to rebuild the device tree online. Examples are provided in the ["Converting port number to area ID"](#) on page 222 section of this chapter.
 - Reboot the device to rebuild the device tree. Some operating systems require a special command to do this, for example `boot -r` in Solaris.
9. For devices that do not bind by PID or have had their PID binding updated, mark online or reassociate the disk devices with the multipathing software and resume I/O over the updated fabric.
 10. Repeat this procedure with the other fabrics.

Offline update

The following steps are intended to provide SAN administrators a starting point for creating site-specific procedures.

1. Schedule an outage for all devices attached to the fabric.
2. Back up all data and verify backups.
3. Shut down all hosts and storage devices attached to the fabric.
4. Disable all switches in the fabric.
5. Change the PID format on each switch in the fabric.
6. Reenable the switches in the updated fabric one at a time.
In a core/edge network, enable the core switches first.
7. After the fabric has reconverged, use the `cfgEnable` command to update zoning.
8. Bring the devices online in the order appropriate to the SAN.
This usually involves starting up the storage arrays first and the hosts last.
9. For any devices manually bound by PID, bring the devices back online, but do not start applications. Update their bindings and reboot again if necessary.
This might involve changing them to the new PIDs, or might (preferably) involve changing to WWN binding.
10. For any devices bound by PID, reboot the device to rebuild the device tree (some operating systems require a special command to do this, such as `boot -r` in Solaris).
11. For devices that do not bind by PID or have had their PID binding updated, bring them back up and resume I/O.
12. Verify that all I/O has resumed correctly.

Hybrid update

It is possible to combine the online and offline methods for fabrics where only a few devices bind by PID. Because any hybrid procedure is extremely customized, it is necessary to work closely with the SAN service provider in these cases.

Changing to Core PID format

In Fabric OS release 4.2.0 and later, Native PID format is not supported; the default format is the Core PID format.

In Fabric OS 3.1.2 and later, Core PID format is the default configuration.

In Fabric OS 2.6.2 and later, Native PID format is the default configuration.

Although the PID format is listed in the configuration file, do not edit the file to change the setting there. Instead, use the CLI `configure` command. When you use the `configure` command, switch databases that contain PID-sensitive information are updated. If you change the setting in the configuration file and then download the edited file, the PID format is changed, but the database entries is not changed, and so they are incorrect.

The following information maps the PID format names to the names used in the management interfaces.

PID format name	Management interface name
Native PID	Switch PID address mode 0
Core PID	Switch PID address mode 1
Extended Edge PID	Switch PID address mode 2

Before changing the PID format, determine whether host reboots are necessary. The section "[Host reboots](#)" on page 214 summarizes the situations that might require a reboot. For example:

```

switch:admin> switchdisable
switch:admin> configure
Configure...

Fabric parameters (yes, y, no, n): [no] y

Domain: (1..239) [1]
BB credit: (1..27) [16]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
WAN_TOV: (1000..120000) [0]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..1) [0]
Switch PID Address Mode: (0..2) [1] < Set mode number here.
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]

```

Changing to Extended Edge PID format

In rare cases, you might be affected by the presence in the fabric of drivers that rely on static binding to the dynamically assigned PID; for example, you might be installing a switch running Fabric OS 4.2.0 into a fabric consisting solely of Fabric OS 2.6.2 and later and Fabric OS 3.1.2 and later switches. In these cases, if you absolutely cannot reboot the affected servers when you upgrade your switches, you can choose Extended Edge PID format. It uses the same PID mapping for the first 16 ports and can support switches and directors with higher port counts. However, because Extended Edge format supports only 128 ports per domain, its use can lead to port addressing issues in directors.

Use the following procedure only if your fabric contains devices that are bound statically and you cannot reboot the host.

1. Determine whether the current switch firmware versions meet the minimum supported version levels.

[Table 47](#) lists the earliest Fabric OS version levels that support Extended Edge PID format. Use this table to determine whether you need to upgrade the firmware in the switches in your fabric before you change the PID format. HP recommends that you download the latest firmware; to download firmware, see ["Obtaining and unzipping firmware"](#) on page 76.

Table 47 Earliest Fabric OS versions for Extended Edge PID format

1 GB Switches	SAN Switch 2/8-EL and SAN Switch 2/16	SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Core Switch 2/64, SAN Director 2/128	Brocade 4Gb SAN Switch for HP p-Class BladeSystem	4/8 SAN Switch, 4/16 SAN Switch, 4/256 SAN Director
2.62	3.12	4.20	5.0.1	5.0.1

2. Update switch firmware as necessary:
 - a. Use the `fabricShow` command to verify the total number of switches in the fabric.
 - b. Download the correct firmware version to each switch as necessary.
 - c. Reboot all switches.
 - d. Verify that the switches form a single fabric and that all domain IDs do not change after forming the fabric.
 - e. Verify that the number of switches is the same.
3. Disable the switch by entering the `switchDisable` command.

4. Change the switch configuration in the fabric to Extended Edge PID format:
 - a. Configure Extended Edge PID (Format 2) on each switch. (See “[PID format changes](#)” on page 224 for a sample `configure` command on an HP StorageWorks switch running Fabric OS 3.1.2 and for a sample `configure` command on an HP StorageWorks switch running Fabric OS 4.2.0 and later.)
 - b. Run the `switchEnable` command all switches.
 - c. Verify that all the switches form a fabric.
 - d. Use the `switchShow` command to verify the ISLs are correct and the device links are correct.
 - e. Use the `fabricShow` command to verify that the number of switches is the same as when you started this procedure.
 - f. Use the `nsAllShow` command to verify the total number of devices is the same as when you started this procedure.
 - g. For dual fabrics, repeat [step 1](#) through [step 4](#) for the other fabric.

The following is an example of the `configure` command on a switch running Fabric OS 3.1.2:

```
Configure...

Fabric parameters (yes, y, no, n): [no] yes

Domain: (1..239) [217]
BB credit: (1..27) [16]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..10) [0]
Switch PID Format : (0..2) [0] 2
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]

Virtual Channel parameters (yes, y, no, n): [no] ^D
Committing configuration...done.
0x102fd500 (tshell): Apr 15 16:53:31
WARNING CONFIG-PIDCHANGE_DISPLACE, 3, Switch PID format changed to Extended
Edge PID Format
```

Example of the configure command on a switch running Fabric OS 5.0.1:

```
configure

Configure...

Fabric parameters (yes, y, no, n): [no] y

Domain: (1..239) [11]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
WAN_TOV: (0..30000) [0]
MAX_HOPS: (7..19) [7]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
Switch PID Format: (1..2) [1] 2
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..27) [16]

Insistent Domain ID Mode (yes, y, no, n): [no]
Virtual Channel parameters (yes, y, no, n): [no]
Zoning Operation parameters (yes, y, no, n): [no]
RSCN Transmission Mode (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
Portlog events enable (yes, y, no, n): [no]
ssl attributes (yes, y, no, n): [no]
http attributes (yes, y, no, n): [no]
snmp attributes (yes, y, no, n): [no]
rpcd attributes (yes, y, no, n): [no]
cfgload attributes (yes, y, no, n): [no]
webtools attributes (yes, y, no, n): [no]

Switch PID format has changed to PID Format 2 ('Extended Edge PID')
```

Converting port number to area ID

Except for the following cases, the area ID is equal to the port number:

- When you perform a port swap operation
- When you enable Extended Edge (also known as *displaced PID*) PID on the switches and directors

If you are using Extended Edge PID format (for example, the 4/256 SAN Director with configuration option 5) and would like to map the output of the port number to the area ID, use the following formula (for ports 0–127):

$$a = (p + 16) \% 128$$

where a is the area, p is the port number, and $\%$ is the remainder. Note that $0 \leq p < 128$.

When the port number is greater than or equal to 128, the area ID and port number are the same.

Figure 13 shows a 4/256 SAN Director with Extended Edge PID.

PID format changes

There are several routine maintenance procedures which might result in a device receiving a new PID. Examples include, but are not limited to:

- Changing compatibility mode settings
- Changing switch domain IDs
- Merging fabrics
- Relocating devices to new ports or new switches (that is, for add, move, and change type operations)
- Updating the Core PID format
- Using hot spare switch ports to deal with failures

In every case where devices employ static PID binding, any such procedure becomes difficult or impossible to execute without downtime.


In some cases, device drivers allow you to specify static PID binding. In these cases, such devices must be identified and their PID binding should be changed to WWN binding.

The following sections contain a basic procedure that summarizes the steps necessary to perform PID format changes without disrupting the fabric, and it provides special procedures for HP/UX and AIX.

Executing the basic procedure

This process should be executed as part of the overall online or offline update process. However, it can be implemented in a stand-alone manner on a non-production fabric or a switch that has not yet joined a fabric.

1. Ensure that all switches in the fabric are running Fabric OS versions that support the addressing mode. HP recommends that you use 2.6.2 for HP StorageWorks 1 GB switches; 3.1.2 for the SAN Switch 2/8-EL and SAN Switch 2/16; 4.2.0 for Core Switch 2/64 and SAN Director 2/128, as well as the SAN Switch 2/8V, SAN Switch 2/16V, and SAN Switch 2/32; and 5.0.1 for 4/256 SAN Director.

 **NOTE:** All switches running any version of Fabric OS 4.0.0 and later are shipped with the Core Switch PID Format enabled, so it is not necessary to perform the PID format change on these switches.

2. Telnet into one of the switches in the fabric.
3. Disable the switch by issuing the `switchDisable` command.
4. Issue the `configure` command.
The configure prompts display sequentially.
5. Enter `y` after the `Fabric parameters` prompt.
6. Enter `1` at the `Core Switch PID Format` prompt.
7. Complete the remaining prompts or press **Ctrl-d** to accept the remaining settings without completing all the prompts.
8. Repeat [step 2](#) through [step 7](#) for the remaining switches in the fabric.
9. Reenable the switch by issuing the `switchEnable` command.

For example:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [1]
R_A_TOV: (4000..120000) [10000]
E_D_TOV: (1000..5000) [2000]
Data field size: (256..2112) [2112]
Sequence Level Switching: (0..1) [0]
Disable Device Probing: (0..1) [0]
Suppress Class F Traffic: (0..1) [0]
SYNC IO mode: (0..1) [0]
Core Switch PID Format: (0..2) [0] 1
Per-frame Route Priority: (0..1) [0]
Long Distance Fabric: (0..1) [0]
BB credit: (1..27) [16]
```

10. After all switches are updated to use the new PID format and reenabled, verify that the fabric has fully reconverged (each switch sees the other switches).
11. Issue the command `cfgEnable [active_zoning_config]` on one of the switches in the fabric to update zoning to use the new PID form.
This does not change the definition of zones in the fabric, but merely causes the lowest level tables in the zoning database to be updated with the new PID format setting. It is necessary to do this only once per fabric; the zoning update propagates to all switches.
At this point, all switches in the fabric are operating in the new addressing mode.

Executing the HP-UX procedure

This procedure is not intended to be comprehensive. It provides a starting point from which a SAN administrator could develop a site-specific procedure for a device that binds by PID, and cannot be rebooted due to uptime requirements.

1. Backup all data and verify the backups.
2. If you are not using multipathing software, stop all I/O going to all volumes connected through the switch/fabric to be updated.
3. If you are not using multipathing software, unmount the volumes from their mount points using the `umount` command.

The proper syntax is `umount mount_point`. For example:

```
umount /mnt/jbod
```

4. If you are using multipathing software, use that software to remove one fabric's devices from its configuration.
5. Deactivate the appropriate volume groups using `vgchange`.

The proper syntax is `vgchange -a n path_to_volume_group`. For example:

```
vgchange -a n /dev/jbod
```

6. Make a backup copy of the volume group directory using `tar` from within `/dev`.

For example:

```
tar -cf /tmp/jbod.tar jbod
```

7. Export the volume group using `vgexport`.

The proper syntax is `vgexport -m mapfile path_to_volume_group`. For example:

```
vgexport -m /tmp/jbod_map /dev/jbod
```

8. Connect to each switch in the fabric.
9. Issue the `switchDisable` command.
10. Issue the `configure` command and change the Core PID format to 1.

11. Issue the command `cfgEnable [effective_zone_configuration]`.

For example:

```
cfgEnable my_zones
```

12. Issue the `switchEnable` command. Enable the core switches first, and then the edges.

13. Clean the `lvmtab` file by using the command `vgscan`.

14. Change to `/dev` and untar the file that was tared in [step 4](#).

For example:

```
tar -xf /tmp/jbod.tar
```

15. Import the volume groups using `vgimport`.

The proper syntax is `vgimport -m mapfile path_to_volume_group physical_volume_path`. For example:

```
vgimport -m /tmp/jbod_map /dev/jbod /dev/dsk/c64t8d0 /dev/dsk/c64t9d0
```

16. Activate the volume groups using `vgchange`.

The proper syntax is `vgchange -a y path_to_volume_group`. For example:

```
vgexport -a y /dev/jbod
```

17. If you are not using multipathing software, mount all devices again and restart I/O.

For example:

```
mount /mnt/jbod
```

18. If you are using multipathing software, reenable the affected path.

The preceding steps do not clean up the results from `ioscan`. When viewing the output of `ioscan`, notice that the original entry is still there, but now has a status of `NO_HW`. For example:

```
# ioscan -funC disk
Class      I    H/W Path                      Driver S/W State   H/W Type           Description
-----
disk       0    0/0/1/1.2.0                      adisk CLAIMED      DEVICE              SEAGATE ST39204LC
                        /dev/dsk/clt2d0 /dev/rdisk/clt2d0
disk       1    0/0/2/1.2.0                      adisk CLAIMED      DEVICE              HP        DVD-ROM 304
                        /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk      319  0/4/0/0.1.2.255.14.8.0          adisk CLAIMED      DEVICE              SEAGATE ST336605FC
                        /dev/dsk/c64t8d0 /dev/rdisk/c64t8d0
disk      320  0/4/0/0.1.18.255.14.8.0         adisk NO_HW         DEVICE              SEAGATE ST336605FC
                        /dev/dsk/c65t8d0 /dev/rdisk/c65t8d0
```

19. To remove the original (outdated) entry, use the command `rmsf` (remove special file).

The proper syntax for this command is `rmsf -a -v path_to_device`. For example:

```
rmsf -a -v /dev/dsk/c65t8d0
```

20. Validate that the entry has been removed by issuing the command `ioscan -funC disk`.

In this example, the `NO_HW` entry is no longer listed:

```
het46 (HP-50001)> ioscan -funC disk
Class      I    H/W Path                      Driver S/W State   H/W Type           Description
-----
disk       0    0/0/1/1.2.0                      adisk CLAIMED      DEVICE              SEAGATE ST39204LC
                        /dev/dsk/clt2d0 /dev/rdisk/clt2d0
disk       1    0/0/2/1.2.0                      adisk CLAIMED      DEVICE              HP        DVD-ROM 304
                        /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk      319  0/4/0/0.1.2.255.14.8.0          adisk CLAIMED      DEVICE              SEAGATE ST336605FC
                        /dev/dsk/c64t8d0 /dev/rdisk/c64t8d0
```

21. Repeat this procedure for all fabrics.

22. Issue the `switchEnable` command. Enable the core switches first, and then the edges.

Executing the AIX procedure

This procedure is not intended to be comprehensive. It provides a starting point from which a SAN administrator can develop a site-specific procedure for a device that binds by PID and cannot be rebooted due to uptime requirements.

1. Backup all data and verify the backups.
2. If you are not using multipathing software, stop all I/O going to all volumes connected through the switch or fabric to be updated.
3. If you are not using multipathing software, issue the `varyoff` command for the volume groups.
The syntax is `varyoffvg volume_group_name`. For example:

```
varyoffvg datavg
```
4. If you are not using multipathing software, unmount the volumes from their mount points using `umount`.
The syntax is `umount mount_point`. For example:

```
umount /mnt/jbod
```
5. If you are using multipathing software, use that software to remove one fabric's devices from its configuration.
6. Remove the device entries for the fabric you are migrating.
For example, if the HBA for that fabric is `fcs0`, issue the command:

```
rmdev -Rdl fcs0
```
7. Connect to each switch in the fabric.
8. Issue the `switchDisable` command.
9. Issue the `configure` command and change the Core PID format to 1.
10. Issue the `configEnable [effective_zone_configuration]` command.
For example:

```
configenable my_config
```
11. Issue the `switchEnable` command. Enable the core switches first, and then the edges.
12. Rebuild the device entries for the affected fabric using the `cfgmgr` command.
For example:

```
cfgmgr -v
```

This command might take several minutes to complete.
13. If you are not using multipathing software, issue the `varyonvg` command for the disk volume groups.
The proper syntax is `varyonvg volume_group_name`. For example:

```
varyonvg datavg
```
14. If you are not using multipathing software, mount all devices again and restart I/O.
For example:

```
mount /mnt/jbod
```
15. If you are using multipathing software, reenabte the affected path.
16. Repeat this procedure for all fabrics.

Swapping port area IDs

If a device that uses port binding is connected to a port that fails, you can use port swapping to make another physical port use the same PID as the failed port. The device can then be plugged into the new port without rebooting the device.

Use the following procedure to swap the port area IDs of two physical switch ports. In order to swap port area IDs, the port swap feature must be enabled, and both switch ports must be disabled. The swapped area IDs for the two ports remain persistent across reboots, power cycles, and failovers.

1. Connect to the switch and log in as admin.

2. Enable the port swap feature:

```
portswapenable
```

3. For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32, issue the following commands:

```
portdisable port1
```

```
portdisable port2
```

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, issue the following commands:

```
portdisable slot/port1
```

```
portdisable slot/port2
```

4. For the 4/8 SAN Switch, 4/16 SAN Switch, SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, Brocade 4Gb SAN Switch for HP p-Class BladeSystem, and SAN Switch 4/32, issue the following command:

```
portswap port1 port2
```

For the Core Switch 2/64, SAN Director 2/128, and 4/256 SAN Director, issue the following command:

```
portswap slot1/port1 slot2/port2
```

5. Verify that the port area IDs have been swapped:

```
portswapshow
```

A table is shows the physical port numbers and the logical area IDs for any swapped ports.

6. Disable the port swap feature:

```
portswapdisable
```

B Configuring interoperability mode

For supported interoperability configurations and restrictions, see the *HP StorageWorks SAN design reference guide* (AA-RMPNW-TE): <http://www.hp.com/go/SANDesignGuide>.

C Using the HP Remote Switch feature

The HP Remote Switch feature (Remote Switch), which aids in ensuring gateway compatibility, was formerly a licensed feature. Its functionality is now available as part of the Fabric OS standard feature set through the use of the `portCfgIslMode` command, which is described in ["Linking through a gateway"](#) on page 33.

Remote Switch, enables you to connect two remote HP StorageWorks switches over an IP network, enabling communication of IP or ATM protocols as well as Fibre Channel traffic.

The Remote Switch functions with the aid of a bridging device or Fibre Channel gateway. The gateway supports both a Fibre Channel physical interface and a secondary, non-Fibre Channel physical interface, such as IP, SONET, or ATM. Remote Switch functions over E_Port connections. With Remote Switch on both fabrics, the gateway accepts Fibre Channel frames from one fabric, tunnels them across the network, and passes them to the other fabric. From the viewpoint of the connected hosts and storage devices, fabrics using Remote Switch interact the same as locally connected switches.

Remote Switch provides many of the same capabilities of normal ISL links including the following:

- Coordinated fabric services: The Remote Switch fabric configuration fully supports all fabric services, including distributed name service, RSCN, and alias service.
- Distributed management: Management tools, such as Advanced Web Tools, Fabric OS, and SNMP are available from both the local switch and the remote switch. Switch management is routed through the Fibre Channel connection; thus, no additional network connection is required between sites.
- Support for ISLs: Sites requiring redundant configurations can connect multiple E_Ports to remote sites by using multiple gateways. Standard Fabric OS routing facilities maximize throughput and provide failover during interruption on the wide area network (WAN) connection.

The Remote Switch feature operates with a gateway. The gateway provides an E_Port interface that links to the E_Port. After the link between the two E_Ports has been negotiated, the gateway E_Port moves to passthrough mode and passes Fibre Channel traffic from the E_Port to the WAN.

The gateway accepts Fibre Channel frames from one side of a Remote Switch fabric, transfers them across a WAN, and passes them to the other side of the Remote Switch fabric.


Remote Switch can be used for these types of gateway devices:

- Fibre Channel over ATM
- Fibre Channel over IP
- Fibre Channel over SONET
- Fibre Channel over DWDM

Most of these gateway devices have enough buffers to cover data transfer over a WAN. The HP StorageWorks switches on each side of the gateway must have identical configurations. Only qualified SFPs should be used.

You must connect the fabrics through the gateway device, and make sure that the `configure` parameters are compatible with the gateway device.

You might be required to reconfigure the following parameters, depending on the gateway requirements:

 **NOTE:** Consult your gateway vendor for supported and qualified configurations.

- R_A_TOV: Specify a Resource Allocation Timeout Value compatible with your gateway device.
- E_D_TOV: Specify an Error Detect Timeout Value compatible with your gateway device
- Data field size: Specify the maximum Fibre Channel data field reported by the fabric. Verify the maximum data field size the network-bridge can handle. Some bridges might not be able to handle a maximum data field size of 2112.
- BB credit: Specify the number of Buffer-to-Buffer credits for Nx_Port devices.

- **Suppress Class F Traffic:** Use this parameter to disable class F traffic. Some network-bridge devices might not have a provision for handling class F frames. In this case, the transmission of class F frames must be suppressed throughout the entire Remote Switch fabric.

To set the access and reconfigure these parameters:

1. Connect to the switch and log in as admin.
 2. Issue the `switchDisable` command to disable the switch.
 3. Issue the `configure` command.
 4. At the Fabric Parameters prompt, enter `yes`.
 5. Press **Enter** to scroll through the Fabric Parameters without changing their values, until you reach the parameter you want to modify.
 6. Specify a new parameter value that is compatible with your gateway device.
 7. Press **Enter** to scroll through the remainder of the configuration parameters. Make sure that the configuration changes are committed to the switch.
 8. Repeat this procedure for all switches in the fabrics to be connected through a gateway device.
- These parameters must be identical on each switch in the fabric and between fabrics connected through the gateway device.

The following example shows how to modify the data field size and suppress class F traffic on a switch:

```
switch:admin> switchdisable
switch:admin> configure
Configure...
Fabric parameters (yes, y, no, n): [no] yes
Domain: (1..239) [3]
  R_A_TOV: (4000..120000) [10000]
  E_D_TOV: (1000..5000) [2000]
  Data field size: (256..2112) [2112] 1000
  Sequence Level Switching: (0..1) [0]
  Disable Device Probing: (0..1) [0]
  Suppress Class F Traffic: (0..1) [0] 1
  VC Encoded Address Mode: (0..1) [0]
  Per-frame Route Priority: (0..1) [0]
  Long Distance Fabric: (0..1) [0]
  BB credit: (1..16) [16]
Virtual Channel parameters (yes, y, no, n): [no]
  Zoning Operation parameters (yes, y, no, n): [no]
  RSCN Transmission Mode (yes, y, no, n): [no]
  NS Operation Parameters (yes, y, no, n): [no]
  Arbitrated Loop parameters (yes, y, no, n): [no]
  System services (yes, y, no, n): [no]
  Portlog events enable (yes, y, no, n): [no]
Committing configuration...done.
switch:admin>
```

D Understanding legacy password behavior

The following sections provide password information for early versions of Fabric OS firmware.

Password management information

Table 48 describes the password standards and behaviors between various versions of firmware.

Table 48 Account and password characteristics matrix

Characteristic	Version 4.0.0	Versions 4.1.0 to 4.2.0	Versions 4.4.0 to 5.0.1
Number of default accounts on the switch	4, chassis-based	Core Switch 2/64—8 for the chassis, 4 per switch SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, SAN Director 2/128—4	Core Switch 2/64—8 for the chassis, 4 per switch SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32, and SAN Director 2/128—4
Default account names	root, factory, admin, user	root, factory, admin, user	root, factory, admin, user
Maximum and minimum amount of characters for a password	0–8 (Standard UNIX)	8–40 characters with printable ASCII	8–40 characters with printable ASCII
Can different switch instances use a different password for the same account login level? For example, can the password for admin for switch 0 be different from the password for admin for switch 1?	No	Yes for Core Switch 2/64. N/A for all other switches.	Yes for Core Switch 2/64. N/A for all other switches.
Does the root account use restricted shell?	No	No	No
When connecting to a factory-installed switch, do you use the default passwords?	Yes	Yes	Yes
Does a user need to know the old passwords when changing passwords using the <code>passwd</code> command?	Yes, except when the root user changes another user's password. This is standard UNIX behavior; Fabric OS does not enforce any additional security.	The old password is required only when changing the password for the same level of user password. Changing the password for a lower-level user does not require the old password. For example, when users connect as admin, the old admin password is required to change the admin password, but the old user password is not required to change the user password.	The old password is required only when changing the password for the same level of user password. Changing the password for a lower-level user does not require the old password. For example, when users connect as admin, the old admin password is required to change the admin password, but the old user password is not required to change the user password.

Table 48 Account and password characteristics matrix (continued)

Characteristic	Version 4.0.0	Versions 4.1.0 to 4.2.0	Versions 4.4.0 to 5.0.1
Can passwd change higher-level passwords? For example, can admin change root password?	Yes, but you must supply the old password of the higher-level account (example <code>root</code>).	Yes; if users connect as admin, they can change the root, factory, and admin passwords. However, if one connects as user, one can change only the user password.	Yes, if users connect as admin, they can change the root, factory and admin passwords after first entering the old password for the respective account.
Can API change passwords?	Yes, only for admin.	Yes, only for admin.	Yes, only for admin.
Can Advanced Web Tools change passwords?	No	No	No
Can SNMP change passwords?	No	No	No

Password prompting behaviors

Table 49 describes the expected password prompting behaviors of various Fabric OS versions.

Table 49 Password prompting matrix

Issue	Version 4.0.0	Version 4.1.0 and later
Must all password prompts be completed for any change to take effect?	No. Partial changes of all four passwords are allowed.	No. Partial changes of all four passwords are allowed.
When does the password prompt appear?	When users connect as root, factory, or admin, the accounts with default password are prompted for change. The accounts with non-default password are not prompted.	When users connect as root, factory, or admin, the accounts with default password are prompted for change. The accounts with non-default password are not prompted.
Is a user forced to answer password prompts before getting access to the firmware?	No, users can press Ctrl-c to get out of password prompting.	No, users can press Ctrl-c to get out of password prompting.
Do users need to know the old root password when answering prompting?	Yes in 4.0.0 *No in 4.0.2 only [*Users do not need to know the old root password if they are running Fabric OS 4.0.2 only.].	No
Are new passwords forced to be set to something different than the old passwords?	Yes	Yes
Is password prompting disabled when security mode is enabled?	Yes	Yes
Is the passwd command disabled until the user has answered password prompting?	No	Yes

Table 49 Password prompting matrix (continued)

Issue	Version 4.0.0	Version 4.1.0 and later
Does password prompting reappear when passwords are changed back to the defaults using the <code>passwd</code> command?	Yes	No
Does password prompting reappear when passwords are changed back to the defaults using the <code>passwdDefault</code> command?	Yes	Yes

Password migration during firmware changes

Table 50 describes the expected outcome of password settings when upgrading or downgrading firmware for various Fabric OS versions.

Table 50 Password migration behavior during firmware upgrade and downgrade

Issue	Versions 4.2.0 to 5.0.1	Versions 4.4.0 to 5.0.1
Which passwords are used when upgrading to a later firmware release for the first time?	Default accounts and passwords are preserved.	Default accounts and passwords are preserved.
Are passwords preserved during subsequent firmware upgrades?	MUA accounts created during a previous upgrade to 4.2.0 are restored.	MUA accounts created during a previous upgrade to 5.0 are restored.
Which passwords are used if downgrading to an earlier firmware for the first time?	Downgrades to 4.2.0 preserve existing default accounts and restore any disabled default accounts. The default account passwords existing immediately prior to the downgrade are preserved. MUA accounts are disabled.	Downgrades to 4.4.0 preserve all existing default accounts, MUA accounts, and passwords. MUA accounts with the <code>switchAdmin</code> role have the same permissions as the user role.
When downgrading to an earlier firmware at subsequent times, which passwords are used?	Downgrades to 4.2.0 preserve existing default accounts and restore any disabled default accounts. The default account passwords existing immediately prior to the downgrade are preserved. MUA accounts are disabled.	Downgrades to 4.4.0 preserve all existing default accounts, MUA accounts, and passwords. MUA accounts with the <code>switchAdmin</code> role have the same permissions as the user role.
When downgrading and then upgrading, what passwords are used?	Any MUA accounts created during the original upgrade are restored. All passwords are unchanged.	All default and MUA accounts and passwords remain unchanged.

Password recovery options

Table 51 describes the options available when one or more types of passwords are lost.

Table 51 Password recovery options

Issue	Version 4.0.0	Versions 4.1.0 and later
If all the passwords are forgotten, what is the password recovery mechanism? Are these procedures non-disruptive recovery procedures?	Contact your switch service provider. A non-disruptive procedure is available.	Contact your switch service provider. A non-disruptive procedure is available.
If a user has only the root password, what is the password recovery mechanism?	Root can change any password by issuing the <code>passwd</code> command.	Use the <code>passwd</code> command to set other passwords. Use the <code>passwdDefault</code> command to set all passwords to their default.
How can boot PROM passwords be recovered?	Not applicable.	Contact your switch service provider and provide the recovery string. See "Setting the boot PROM password" on page 67 for instructions on setting the password with a recovery string.
How can a user, admin, or factory password be recovered?	See "Recovering forgotten passwords" on page 71.	

E Zone merging scenarios

Table 52 provides information on merging zones and the expected results.

Table 52 Zone merging scenarios


Description	Switch A	Switch B	Expected Results
Switch A has a defined configuration. Switch B does not have a defined configuration.	defined: none cfg1: none zone1: ali1; ali2 effective: none	defined: none effective: none	Configuration from Switch A to propagate throughout the fabric in an inactive state, because the configuration is not enabled.
Switch A has a defined and enabled configuration. Switch B has a defined configuration but no effective configuration.	defined: cfg1 zone1: ali1; ali2 effective: cfg1	defined: cfg1 zone1: ali1; ali2 effective: none	Configuration from Switch A to propagate throughout the fabric. The configuration is enabled after the merge in the fabric.
Switch A and Switch B have the same defined configuration. Neither have an enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	No change (clean merge).
Switch A and Switch B have the same defined and enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: cfg1	defined: cfg1 zone1: ali1; ali2 effective: cfg1	No change (clean merge).
Switch A does not have a defined configuration. Switch B has a defined configuration.	defined: none effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	Switch A absorbs the configuration from the fabric.
Switch A does not have a defined configuration. Switch B has a defined configuration.	defined: none effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1	Switch A absorbs the configuration from the fabric, with cfg1 as the effective cfg.
Switch A and Switch B have the same defined configuration. Only Switch B has an enabled configuration.	defined: cfg1 zone1: ali1; ali2 effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1	Clean merge with cfg1 as the effective cfg.
Switch A and Switch B have different defined configurations. Neither have an enabled zone configuration.	defined: cfg2 zone2: ali3; ali4 effective: none	defined: cfg1 zone1: ali1; ali2 effective: none	Clean merge. The new cfg is a composite of the two. defined: cfg1 zone1: ali1; ali2 cfg2: zone2: ali3; ali4 effective: none
Switch A and Switch B have different defined configurations. Switch B has an enabled configuration.	defined: cfg2 zone2: ali3; ali4 effective: none	defined: cfg1 zone1: ali1; ali2 effective: cfg1	Clean merge. The new cfg is a composite of the two, with cfg1 as the effective cfg.
There is an effective cfg mismatch.	defined: cfg1 zone1: ali1; ali2 effective: cfg1 zone1: ali1; ali2	defined: cfg2 zone2: ali3; ali4 effective: cfg2 zone2: ali3; ali4	Fabric segments due to a zone conflict cfg mismatch.

Table 52 Zone merging scenarios (continued)

Description	Switch A	Switch B	Expected Results
There is a cfg content mismatch.	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone1: ali3; ali4 effective: irrelevant	Fabric segments due to a zone conflict content mismatch.
	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone1: ali1; ali4 effective: irrelevant	Fabric segments due to a zone conflict content mismatch.
Same content, with a different effective cfg name.	defined: cfg1 zone1: ali1; ali2 effective: cfg1 zone1: ali1; ali2	defined:cfg2 zone1: ali1; ali2 effective:cfg2 zone1: ali1; ali2	Fabric segments due to a zone conflict cfg mismatch.
Same content with different zone name.	defined: cfg1 zone1: ali1; ali2 effective: irrelevant	defined: cfg1 zone2: ali1; ali2 effective: irrelevant	Fabric segments due to a zone conflict content mismatch.
Same content with a different alias name.	defined: cfg1 ali1: A; B effective: irrelevant	defined:cfg1:ali2: A; B effective: irrelevant	Fabric segments due to a zone conflict content mismatch.
Same name with different types.	effective: zone1: MARKETING	effective: cfg1: MARKETING	Fabric segments due to a zone conflict type mismatch.
Same name with different types.	effective: zone1: MARKETING	effective: alias1: MARKETING	Fabric segments due to a zone conflict type mismatch.
Same name with different types.	effective:cfg1: MARKETING	effective:alias1: MARKETING	Fabric segments due to a zone conflict type mismatch.


F Upgrading firmware in single-CP mode

For all HP StorageWorks switches and directors, the `firmwareDownload` command, by default, performs a full installation, automatic reboot (autoreboot), and automatic firmware commit (autocommit). Automatic reboot and automatic commit modes are not selectable by default; however, they become selectable when single-CP mode is enabled by entering the `-s` option on the command line. In this case, `firmwareDownload` disables autoreboot and continues to enable autocommit mode by default.

 **NOTE:** Use the following procedures only if instructed to do so by your service provider.

Your service provider might ask you to perform this procedure on your HP StorageWorks switch—or on one or both CPs in director models—under the following circumstances:

- To prevent the firmware commit that occurs after downloading, so that you can restore previous versions
- To control the timing of the execution of the `haReboot` command, so that you can prestage the firmware ahead of time

 **NOTE:** Make sure that your switch is in a steady state (for example, no fabric configuration or ISL/cable changes) before issuing the `haReboot` command. The `haReboot` command synchronizes firmware versions and associated data in real time. If your switch is not in a steady state when you issue the `haReboot` command, the switch performs a cold restart for the affected area, which can disrupt the ASIC and machine traffic.

Your service provider may ask you to perform the single-CP blade procedure on HP StorageWorks SAN Directors if a CP blade fails and the replacement CP blade is running a version of firmware that cannot synchronize with the current active CP blade.

For information about messages that might appear during the procedures, see the *HP StorageWorks Fabric OS 5.x diagnostics and system error messages reference guide*.

Upgrading HP StorageWorks SAN Switch 2/8V, SAN Switch 2/16V, SAN Switch 2/32, and SAN Switch 4/32

Specify the `-s` option on the command line for single-CP mode. You are prompted for other options and can enable autocommit and autoreboot using the following procedure.

1. Connect to the switch and log in as admin.
2. Issue the following command:

```
firmwaredownload -s
```
3. Enter the IP address of the FTP server where the firmware is stored.
4. Enter your user name for the server.
5. Enter the full path to the firmware file on the server.

For example:

```
/pub/v5.0.1/release.plist
```

6. Enter your password.

7. Answer the next prompts as indicated:

```
Do Auto Commit after reboot [Y]: y
```

If you specify no, you must manually issue the `firmwareCommit` command.

```
Reboot system after download [N]: y
```

The default is no. If you take the default, you must later use the `haReboot` command to perform an HA reboot manually.

In Fabric OS 4.4.0 or later, the Full Install option is not available.

8. Wait for the firmware download to finish.

9. Start a new telnet session and use the `firmwareDownloadStatus` command to check the status.

For example:

```
switch: admin> firmwaredownload -s
Server Name or IP Address: 192.1.2.3
User Name: JohnDoe
File Name: /pub/v5.0.1/release.plist
Password: *****
Do Auto Commit after reboot [Y]: y
Reboot system after download [N]: y
Firmwaredownload has started.
.
.
.
```

Upgrading a single Core Switch 2/64 or SAN Director 2/128 blade

Although it is possible to upgrade firmware on one CP blade at a time, not do so under normal circumstances, because it might disrupt switch operations if it is not executed in the proper sequence, or if the version of firmware is down-level and part of an unsupported version.

When the two CP blades are not running the same firmware versions, it might be necessary to disable one or the other blades to maintain fabric stability. For information on the commands used to achieve this, see the `haDisable` and `haFailover` commands in the *HP StorageWorks Fabric OS 5.x command reference guide*.

The following procedure allows you to upgrade a single CP blade. This procedure can be used with Fabric OS 4.0.0d and later.

1. Connect to the switch and log in as admin.
2. Issue the `haShow` command to determine which CP blade is the active CP and which one is the standby CP. In the following example, the active CP blade is CP0, and the standby CP blade is CP1:

```
switch:admin> hashow
Local CP (Slot 5, CP0): Active
Remote CP (Slot 6, CP1): Standby, Healthy
HA enabled, Heartbeat Up, HA State in sync
```

3. Log in to the standby CP blade as admin.
4. Issue the following command to upgrade a new version of the firmware to the standby CP blade:

```
firmwaredownload -s
```

5. Enter the IP address of the FTP server where the firmware is stored.
6. Enter your user name for the server.

7. Enter the full path to the firmware file on the server.

For example:

```
/pub/v5.0.1/release.plist
```


8. Enter your password.
9. Answer the next prompts as indicated:

```
Do Auto Commit after reboot [Y]: y
```

If you answer no in the previous example, you must manually issue the `firmwareCommit` command.

```
Reboot system after download [N]: y
```

The default is no. If you take the default, you must later use the `haReboot` command to perform an HA reboot manually.

 **NOTE:** After you upgrade to Fabric OS 4.4.0 or later, the Full Install option is no longer available.

For example:

```
switch: admin> firmwaredownload -s  
Server Name or IP Address: 192.1.2.3  
User Name: JohnDoe  
File Name: /pub/v5.0.1/release.plist  
Password: *****  
Full Install (Otherwise upgrade only) [Y]: y  
Do Auto Commit after reboot [Y]: y  
Reboot system after download [N]: y
```

10. Wait for the firmware download to finish. (Start a new telnet session and use the `firmwareDownloadStatus` command to check the status.)
11. Issue the `haShow` command to verify that the two CP blades are synchronized.
12. If you set the option to reboot to no, reboot the standby CP blade.
13. Log in to the same CP blade and issue the `firmwareDownloadStatus` command to verify that the firmware has downloaded successfully and has either committed or is in the process of doing so. (If you set the option to perform an autoccommit after reboot to no in [step 9](#), you must issue the `firmwareCommit` command manually.)
14. Issue the `haFailover` command to make the current CP blade active (with the updated firmware).
15. Repeat [step 1](#) through [step 13](#) on the other CP blade if the firmware versions are different.

Index

A

- accessing switches and fabrics 42
- account ID 21
- account privilege levels 22
- activating
 - a switch certificate 56
 - ports on demand 30
- adding
 - a new switch or fabric 195
 - and removing FICON CUP licenses 116
 - custom filter-based monitors 205
 - end-to-end monitors 201
 - filter-based monitors 205
 - standard filter-based monitors 205
 - switches to a zone 195
 - zone members 187
- advanced performance monitoring commands 199
- AIX procedure, PID 227
- analyzing connection problems 145
- assigning a static route 98
- audience 13
- authorized reseller, HP 15

B

- backing up
 - a configuration 73
 - and restoring configurations, FICON 116
- basic
 - card management 88
 - PID procedure 224
- beaconing mode 95
- blocking listeners 42
- boot PROM password 67
- browser and Java support 53
- buffer-limited port 163

C

- card management 88
- changes to configuration data 214
- changing
 - an account password 45
 - to core PID format 219
 - to extended edge PID format 220
- chassisshow command 34
- checking
 - connected switches 77
 - status 34
- choosing
 - a CA 54
 - an extended ISL mode 164
- clearing
 - FICON management database 107, 111
 - monitor counters 209

CLI 18

- collecting performance data 211
- combining SilkWorm 12000 and 24000 cards in one chassis 95
- command
 - advanced performance monitoring 199
 - chassisshow 34
 - configupload 76
 - fabricshow 34
 - hashow 34
 - licenseadd 27
 - licenseremove 28
 - licenseshow 28
 - nsallshow 35
 - nsshow 35
 - slotshow 34
 - switchshow 34
- configuration
 - FICON environment switched point-to-point 106
 - FICON environment, cascaded 106
 - high-integrity fabric 109
 - new SilkWorm 24000 with two domains 93
 - save to a host 73
 - settings, FICON environment 108
- configuring a single switch 108
- configuring an extended ISL 165
- configuring for SNMP 58
- configuring for syslogd 138
- configuring for the SSL protocol 52
- configuring secure file copy 67
- configuring the browser 56
- configuring the host 138
- configuring the radius server 46
- configuring the server database 126
- configuring the switch 50
- configuring the switch in a FICON environment 108
- configuring the telnet interface 41
- connecting to devices 31
- connecting to other switches 31
- connecting to the command line interface 21
- connection
 - serial 21
 - telnet 21
- conserving power 89
- controlling access 124
- controlling topology discovery 127
- conventions
 - document 14
 - text symbols 14
- converting an installed SilkWorm 24000 to support two domains 94
- core/edge topology and ISL trunking 168
- correcting device login issues 152
- correcting I2C bus errors 151

- correcting link failures [157](#)
- correcting marginal links [159](#)
- correcting zoning setup issues [149](#)
- CRC errors, displaying [200](#)
- creating a zone [187](#)
- creating and maintaining user-defined accounts [43](#)
- creating and maintaining zones [187](#)
- creating and managing zone aliases [185](#)
- creating and modifying zoning configurations [190](#)
- customizing the chassis name [29](#)
- customizing the switch name [28](#)

D

- database, clearing in a FICON environment [107](#)
- date and time [24](#)
- default names [28](#)
- default password [23](#)
- defined zone configuration [180](#)
- deleting
 - end-to-end monitors [204](#)
 - filter-based monitors [206](#)
- deleting a zone [188](#)
- deleting end-to-end monitors [204](#)
- deleting filter-based monitors [206](#)
- designing fabric for trunking [168](#)
- deskew values
 - displaying [173](#)
- device
 - connecting [31](#)
- device-based routing [97](#)
- disabled zone configuration [180](#)
- disabling and enabling a port [30](#)
- disabling and enabling a switch [29](#)
- disabling and enabling cards [89](#)
- displaying
 - CRC error count [200](#)
 - end-to-end mask [203](#)
 - node identification data, FICON environments [110](#)
 - registered listeners for link incidents, FICON environment [110](#)
- displaying additional help topics [19](#)
- displaying and clearing the CRC error count [200](#)
- displaying and deleting certificates [57](#)
- displaying command help [18](#)
- displaying configuration settings [73](#)
- displaying information [110](#)
- displaying mode register bit settings [113](#)
- displaying monitor counters [207](#)
- displaying the fmsmode setting [113](#)
- displaying trunking information [173](#)
- document
 - conventions [14](#)
 - related documentation [13](#)
- domain ID
 - FICON environment [32](#)
- domain ID, insistent [105](#)
- downloading configurations across a fabric [76](#)
- downloading firmware [78](#)

E

- effective zone configuration [180](#)
- enabling and disabling FICON management server mode [112](#)
- enabling and disabling ISL trunking [171](#)
- enabling and disabling local authentication [52](#)
- enabling and disabling the platform services [123](#)
- encryption [53](#)
- end-to-end monitoring [201](#)
- end-to-end monitors
 - adding [201](#)
 - deleting [204](#)
 - displaying the mask [203](#)
 - restoring configuration [211](#)
 - saving configuration [211](#)
 - setting a mask [202](#)
- ensuring network security [40](#)
- evaluating the fabric [216](#)
- example
 - chassisshow [34](#)
 - fabricshow [34](#)
 - nsallshow [35](#)
 - slotshow [34](#)
- exchange-based routing [97](#)
- extended link buffer allocation [163](#)

F

- fabric
 - high integrity [106](#)
- fabric access [42](#)
- fabric connectivity [34](#)
- fabric considerations [163](#), [168](#)
- fabric, designing for trunking [168](#)
- fabricshow command [34](#)
- fans, status of [135](#)
- feature licenses [26](#)
- FICON environment
 - cascaded configuration [106](#)
 - changing domain id [32](#)
 - configuration settings [108](#)
 - configuring switches [108](#)
 - disabling IDID mode [107](#)
 - displaying
 - link incidents [107](#)
 - registered listeners for link incidents [110](#)
 - enabling IDID mode [107](#)
 - high integrity fabric [106](#)
 - identifying port swapping nodes [111](#)
 - monitoring FRU failures [111](#)
 - node identification data, displaying [110](#)
 - switched point-to-point configuration [106](#)
- filter-based monitors [204](#)
 - adding [205](#)
 - deleting [206](#)
 - restoring configuration [211](#)
 - saving configuration [211](#)
- firmware download [78](#)
- frame transfer with brocade remote switch [231](#)

fru failures 111
fru failures, monitoring in FICON environments 107, 111

G

gateway 231
gateway, remote switch 231
gathering information for technical support 144
generating
 batch of licenses 27
generating a public/private key 54
generating and storing a csr 54

H

hard zoning 181
hardware-enforced zoning 181
hashow command 34
help information 18
help, obtaining 14, 15
high availability (HA) 34
high integrity fabric 106
host reboots 214
host-based zoning 178
HP
 authorized reseller 15
 storage web site 15
 Subscriber's choice web site 15
 technical support 14
HP/UX procedure 225
hybrid update 219

I

ID, account 21
identifying
 ports from the tag field 117
identifying media-related issues 155
Identifying ports
 by slot and port number 87
identifying ports 87
 by port area ID 88
IDID 105
IDID mode
 enabling and disabling in a FICON environment 109
impact of changing the fabric PID format 213
inaccurate information in the system message log 161
initializing trunking on ports 169
insistent domain ID 105
installing a root certificate to the Java plug-in 57
installing a switch certificate 55
intermix mode 105
interswitch link 31
ISL 31

J

Java support, SSL 53

K

key

license, generating on the web 26

L

license key
 activating 27
license keys
 generating 26
licenseadd command 27
licensed features 26
licenseremove command 28
licenseshow command 27
link incidents
 displaying in a FICON environment 107, 110
linking through a gateway 33
listing link characteristics 174
login
 switch 21
long distance ISLs 163
LUN masking 178
LWL
 ISL Trunking support for 167

M

maintaining configurations 73
maintaining firmware 76
maintaining licensed features 26
making basic connections 31
managing zoning configurations in a fabric 193
mask for end-to-end monitors
 displaying 203
 setting 202
monitoring end-to-end performance 201
monitoring filter-based performance 204
monitoring ISL performance 207
monitoring traffic 169
monitoring trunks 207
monitors
 clearing counters 209
most common problem areas 143

N

name server zoning 178
network security 40
node identification data 110
nsallshow command 35
nsshow command 35

O

obtaining and unzipping firmware 76
obtaining certificates 55
obtaining slot information 92
offline update 218
online update 218
optimizing resources through zoning 177

P

password 21
 boot prom 67

- default 23
- password management information 233
- password migration during firmware changes 235
- password prompting behaviors 234
- password recovery options 236
- passwords
 - recovering forgotten passwords 71
- perfaddeemonitor command 201
- perfaddIPmonitor command 205
- perfaddusermonitor command 205
- perfcfgrestore command 211
- perfcfgsave command 211
- perfdleemonitor command 204
- perfdelfiltermonitor command 206
- performance monitoring commands 199
- performing PID format changes 222
- perfsetportteemask command 202
- perfshowpalcrc command 200
- perfshowportteemask command 203
- persistently enabling/disabling ports 115
- PID
 - AIX procedure 227
 - binding 213
- PKI 52
- planning the update procedure 218
- policies, routing 97
- port
 - buffer-limited 163
- port and switch naming standards 116
- port swapping nodes, identifying in FICON environments 111
- port-based routing 97
- ports
 - identifying by port area ID 88
 - identifying by slot and port number 87
 - status of 132
- ports, swapping 111
- powering off a card 88
- powering port cards on and off 88
- preparing a switch 108
- printing hard copies of switch information 76
- privileges in accounts 22
- procedural differences between fixed-port and variable-port switches 17
- public key infrastructure encryption 52

R

- recognizing buffer underallocation 174
- recognizing MQ-WRITE errors 151
- recognizing the port initialization and FCP auto discovery process 161
- recording configuration information 118
- recovering forgotten passwords 71
- recovery password 69
- recovery string, boot PROM password 67
- registered listeners 110
- related documentation 13
- remote switch 231
- removing

- end-to-end monitors 204
- filter-based monitors 206
- removing members
 - zone 188
- resolving zone conflicts 197
- restoring a configuration 74
- restoring a segmented fabric 148
- restoring monitor configuration 211
- restoring the system configuration settings 74
- routing
 - assigning a static route 98
- routing policies 97
- rules for configuring zones 184

S

- saved zone configuration 180
- saving and restoring monitor configurations 211
- saving monitor configuration 211
- scope and references 17
- secure shell (ssh) 40
- security 40
- security and zoning 197
- selecting a PID format 215
- serial connection 21
- setting a mask for end-to-end monitors 202
- setting a unique domain id 109
- setting chassis configurations 90
- setting mask for end-to-end monitors 202
- setting mode register bits 114
- setting port speeds 171
- setting the boot prom password 67
- setting the card beacon mode 95
- setting the date and time 24
- setting the default account passwords 22
- setting the IP address 22
- setting the security level 59
- setting the switch date and time 24
- setting up automatic trace dump transfers 140
- setting up RADIUS AAA service 45
- setup summary 112
- slotshow command 34
- SNMP 58, 59, 63, 64, 65
- SNMPv1 61, 62
- SNMPv3 60
- specifying frame order delivery 98
- specifying the routing policy 97
- splitting a fabric 197
- standard filter-based monitors 205
- standard trunking criteria 168
- static PID mapping errors 214
- static route 98
- storage-based zoning 178
- Subscriber's choice, HP 15
- summary of PID formats 213
- summary of SSL procedures 53
- supportsave command 140
- swapping port area IDs 227
- swapping ports 111
- switch

- configuring in a FICON environment 108
- system status 130
- switch access 42
- switch names 28
- switchshow command 34
- SWL, ISL Trunking support for 167
- symbols in text 14

T

- tag field, interpreting 117
- technical support, HP 14
- telnet connection 21
- temperature, status of 136
- text symbols 14
- time and date 24
- tracking and controlling switch changes 35
- traffic patterns
 - planning for 169
- troubleshooting 116
- troubleshooting certificates 58
- troubleshooting firmware downloads 85
- troubleshooting trunking problems 174
- trunking
 - displaying information 173
 - over distance 166
- trunking over extended fabrics 173

U

- upgrading SilkWorm 3016, 3250, 3850, 3900, and 4100 switches 79
- user-defined filter-based monitors 205
- using dynamic load sharing 99
- using FICON CUP 111
- using legacy commands for SNMPv1 62
- using the snmpconfig command 59
- using zoning to administer security 197

V

- verify
 - device connectivity 31
 - fabric connectivity 34
 - high availability (HA) 34
- viewing
 - fan status 135
 - port status 132
 - power supply status 135
 - temperature status 136
- viewing and saving diagnostic information 140
- viewing equipment status 135
- viewing port information 132
- viewing power-on self test 129
- viewing routing information along a path 102
- viewing routing path information 100
- viewing switch status 130
- viewing the port log 137
- viewing the system message log 136

W

- web sites
 - HP storage 15
 - HP Subscriber's choice 15
- without a recovery string 69
- working with domain IDs 32

Z

- zone
 - adding members 187
 - adding switches 195
 - creating 187
 - deleting 188
 - removing members 188
- zone aliases 180
- zone configurations 180
- zone objects 179
- zone types 178
- zoning
 - administering security 197
- zoning and PDCM considerations 116
- zoning concepts 178
- zoning enforcement 181
- zoning schemes 180
- zoning terminology 177

